The Hidden Convexity of Spectral Clustering

James Voss* vossj@cse.ohio-state.edu
Mikhail Belkin* mbelkin@cse.ohio-state.edu
Luis Rademacher* lrademac@cse.ohio-state.edu

October 5, 2015

arXiv:1403.0667v2 [cs.LG] 1 Oct 2015

Abstract

In recent years, spectral clustering has become a standard method for data analysis used in a broad range of applications. In this paper we propose a new class of algorithms for multiway spectral clustering based on optimization of a certain "contrast function" over the unit sphere. These algorithms, partly inspired by certain Indepenent Component Analysis techniques, are simple, easy to implement and efficient.

Geometrically, the proposed algorithms can be interpreted as hidden basis recovery by means of function optimization. We give a complete characterization of the contrast functions admissible for provable basis recovery. We show how these conditions can be interpreted as a "hidden convexity" of our optimization problem on the sphere; interestingly, we use efficient convex maximization rather than the more common convex minimization. We also show encouraging experimental results on real and simulated data.

keywords: spectral clustering | convex optimization | basis recovery

1 Introduction

Partitioning a dataset into classes based on a similarity between data points, known as cluster analysis, is one of the most basic and practically important problems in data analysis and machine learning. It has a vast array of applications from speech recognition to image analysis to bioinformatics and to data compression. There is an extensive literature on the subject, including a number of different methodologies as well as their various practical and theoretical aspects [12].

In recent years spectral clustering—a class of methods based on the eigenvectors of a certain matrix, typically the graph Laplacian constructed from data—has become a widely used method for cluster analysis. This is due to the simplicity of the algorithm, a number of desirable properties it exhibits and its amenability to theoretical analysis. In its simplest form, spectral bi-partitioning is an attractively straightforward algorithm based on thresholding the second bottom eigenvector of the Laplacian matrix of a graph. However, the more practically significant problem of multiway spectral clustering is considerably more complex. While hierarchical methods based on a sequence of binary splits have been used, the most common approaches use k-means or weighted k-means clustering in the spectral space or related iterative procedures [19, 16, 2, 26]. Typical algorithms for multiway spectral clustering follow a two-step process:

1. *Spectral embedding:* A similarity graph for the data is constructed based on the data's feature representation. If one is looking for k clusters, one constructs the embedding using the bottom k eigenvectors of the graph Laplacian (normalized or unnormalized) corresponding to that graph.

*Department of Computer Science and Engineering, the Ohio State University

2. *Clustering:* In the second step, the embedded data (sometimes rescaled) is clustered, typically using the conventional/spherical k-means algorithms or their variations.

In the first step, the spectral embedding given by the eigenvectors of Laplacian matrices has a number of interpretations. The meaning can be explained by spectral graph theory as relaxations of multiway cut problems [21]. In the extreme case of a similarity graph having k connected components, the embedded vectors reside in $\mathbb{R}^k$, and vectors corresponding to the same connected component are mapped to a single point. There are also connections to other areas of machine learning and mathematics, in particular to the geometry of the underlying space from which the data is sampled [4].

In our paper we propose a new class of algorithms for the second step of multiway spectral clustering. The starting point is that when k clusters are perfectly separate, the spectral embedding using the bottom k eigenvectors has a particularly simple geometric form. For the unnormalized (or asymmetric normalized) Laplacian, it is simply a (weighted) orthogonal basis in k-dimensional space, and recovering the basis vectors is sufficient for cluster identification. This view of spectral clustering as basis recovery is related to previous observations that the spectral embedding generates a discrete weighted simplex (see [22] and also [13] for some applications). For the symmetric normalized Laplacian, the structure is slightly more complex, but is, as it turns out, still suitable for our analysis, and, moreover the algorithms can be used without modification.

The approach taken in our paper relies on an optimization problem resembling certain Independent Component Analysis techniques, such as FastICA (see [11] for a broad overview). Specifically, we show that the problem of identifying k clusters reduces to maximizing a certain "admissible" contrast function over a $(k-1)$-sphere. Each local maximum of such a function on the sphere corresponds to exactly one cluster in the data. The main theoretical contribution of our paper is to provide a complete characterization of the admissible contrast functions for geometric basis recovery. We show that such contrast functions have a certain "hidden convexity" property and that this property is necessary and sufficient for guaranteed recovery[1] (Section 2). Rather than the more usual convex minimization, our analysis is based on *convex maximization* over a (hidden) convex domain. Interestingly, while *maximizing* a convex function over a convex domain is generally difficult (even maximizing a positive definite quadratic form over the continuous cube $[0,1]^n$ is NP-hard[2]), our setting allows for efficient optimization.

Based on this theoretical connection between clusters and local maxima of contrast functions over the sphere, we propose practical algorithms for cluster recovery through function maximization. We discuss the choice of contrast functions and provide running time analysis. We also provide a number of encouraging experimental results on synthetic and real-world datasets.

Finally, we note connections to recent work on geometric recovery. The paper [1] uses the method of moments to recover a continuous simplex given samples from the uniform probability distribution. Like our work, it uses efficient enumeration of local maxima of a function over the sphere. In [10], one of the results shows recovery of parameters in a Gaussian Mixture Model using the moments of order three and can be thought of as a case of the basis recovery problem.

The paper is structured as follows: in Section 2 we state the main theoretical results of the paper providing complete description of allowable contrast functions for weighted basis recovery as well as briefly outlining its connection to spectral clustering. In Sections 3 and 4 we introduce spectral clustering and formulate it in terms of basis learning. In Section 5 we provide the main theoretical results for basis recovery in the spectral clustering setting. Section 6 discusses the

[1] Interestingly, there are no analogous recovery guarantees in the ICA setting except for the special case of cumulant functions as contrasts. In particular, typical versions of FastICA are known to have spurious maxima [23].

[2] This follows from [7] together with Fact 1 below.

algorithms, choices of contrast functions and implementations. Some experimental results are given in Section 7.

2 Summary of the Theoretical Results

In this section we state the main theoretical results of our paper on weighted basis recovery and briefly show how they can be applied to spectral clustering.

A Note on Notation. Before proceeding, we define some notations used throughout the paper. The set $\{1, 2, \ldots, k\}$ is denoted by $[k]$. For a matrix B, b_{ij} indicates the element in its i^{th} row and j^{th} column. The i^{th} row vector of B is denoted $b_{i\cdot}$, and the j^{th} column vector of B is denoted $b_{\cdot j}$. For a vector v, $\|v\|$ denotes its standard Euclidean 2-norm. Given two vectors u and v, $\langle u, v\rangle$ denotes the standard Euclidean inner produce between the vectors. We denote by $\mathbf{1}_{\mathcal{S}}$ the indicator vector for the set $\mathcal{S}$, i.e. the vector which is 1 for indices in $\mathcal{S}$ and 0 otherwise. The null space of a matrix M is denoted $\mathcal{N}(M)$. We denote the unit sphere in $\mathbb{R}^d$ by S^{d-1}. For points $p_1, \ldots, p_m$, $\text{conv}(p_1, \ldots, p_m)$ will denote their convex hull. All angles are given in radians, and $\angle(u, v)$ denotes the angle between the vectors u and v in the domain $[0, \pi]$. Finally, for $\mathcal{X}$ a subspace of $\mathbb{R}^m$, $P_{\mathcal{X}}$ denotes the square matrix corresponding to the orthogonal projection from $\mathbb{R}^m$ to $\mathcal{X}$.

Recovering a Weighted Basis. The main technical results of this paper deal with reconstructing a weighted basis by optimizing a certain contrast function over a unit sphere. We show that for certain functions, their maxima over the sphere correspond to the directions of the basis vectors. We give a complete description for the set of such functions, providing necessary and sufficient conditions.

More formally, consider a set $\{Z_1, \ldots, Z_m\}$ of m orthonormal vectors in $\mathbb{R}^m$. These vectors form a hidden basis of the space. We define a function $F_g : S^{m-1} \to \mathbb{R}$ in terms of a "contrast function" g and strictly positive weights α_i, β_i as follows:

$$F_g(u) := \sum_{i=1}^{m} \alpha_i g(\beta_i |\langle u, Z_i\rangle|) \ . \tag{1}$$

We will provide a complete description of when the directions $Z_1, \ldots, Z_m$ can be recovered from the local maxima of F_g for arbitrary weights α_i, β_i. This process of finding the local maxima of F_g can be thought of as weighted *basis recovery.*

Here and everywhere else in the paper, we consider contrast functions $g : [0, \infty) \to \mathbb{R}$ that are continuous on $[0, \infty)$ and twice continuously differentiable on $(0, \infty)$. It turns out that the desirable class of function can be described by the following properties:

P1. Function $g(\sqrt{x})$ is strictly convex.

P2. The (right) derivative at the origin, $\frac{d}{dx}(g(\sqrt{x}))|_{x=0^+}$, is 0 or $-\infty$.

The main idea underlying the proposed framework for weighted basis recovery comes from property P1. In particular, using just this property, it can be shown that the local maxima of F_g are contained in the set $\{\pm Z_i : i \in [m]\}$. The idea is to perform a change of variable to recast maximization of F_g over the unit sphere as a convex maximization problem defined over a (hidden) convex domain. We sketch the proof in order to illustrate this hidden role of convexity in weighted basis recovery.

Proof sketch: Maxima of F_g are contained in $\{\pm Z_i : i \in [m]\}$.

We will need the following fact about convex maximization [17, Chapter 32].

For a convex set K, a point $x \in K$ is said to be an *extreme point* if x is not equal to a strict convex combination of two other points of K.

Fact 1. *Let $K \subseteq \mathbb{R}^n$ be a convex set. Let $f : K \to \mathbb{R}$ be a strictly convex function. Then the set of local maxima of f on K is contained in the set of extreme points of K.*

As $Z_1, \ldots, Z_m$ form an orthonormal basis of the space, we may simplify notation by working in the hidden coordinate system in which $Z_1, \ldots, Z_m$ are the canonical vectors $e_1, \ldots, e_m$ respectively. Let $\Delta^{m-1} := \text{conv}(e_1, \ldots, e_m)$ denote a (hidden) simplex. We will make use of the change of variable $\psi : S^{m-1} \to \Delta^{m-1}$ defined by $\psi_i(u) = u_i^2$. In particular, we define a family of functions $h_i : [0, \infty) \to \mathbb{R}$ for $i \in [m]$ by $h_i(t) = \alpha_i g(\beta_i \sqrt{t})$, and we define a function $H : \Delta^{m-1} \to \mathbb{R}$ as $H(t) = \sum_{i=1}^m h_i(t_i)$. Using assumption P1, it can be seen that H is a strictly convex function defined on a convex domain. Further, for any $u \in S^{m-1}$, $(H \circ \psi)(u) = F_g(u)$. Using this equality, we see that u is a maximum of F_g if and only if $\psi(u)$ is a maximum of H.

The extreme points of Δ^{m-1} are $Z_1, \ldots, Z_m$. By Fact 1, the maxima of H are contained in the set $\{Z_i : i \in [m]\}$. Hence, the maxima of F_g are contained in $\psi^{-1}(\{Z_i : i \in [m]\}) = \{\pm Z_i : i \in [m]\}$. □

We have demonstrated that F_g has no local maxima outside of the set $\{\pm Z_i : i \in [m]\}$; however, we have not demonstrated that the directions $\{\pm Z_i : i \in [m]\}$ actually are local maxima of F_g. In general, both P1 and P2 are required to guarantee that $\{\pm Z_i : i \in [m]\}$ is a complete enumeration of the local maxima of F_g. More formally, we have the following main theoretical results (proven in Appendix A):

Theorem 2 (Sufficiency)**.** *Let $\alpha_1, \ldots, \alpha_m$ and $\beta_1, \ldots, \beta_m$ be strictly positive constants. Let $g : [0, \infty) \to \mathbb{R}$ be a continuous function which is twice continuously differentiable on $(0, \infty)$ satisfying properties P1 and P2. If $F_g : S^{m-1} \to \mathbb{R}$ is constructed from g according to equation (1), then all local maxima of F_g are contained in the set $\{\pm Z_i\}_{i=1}^m$ of basis vectors. Moreover, each basis vector $\pm Z_i$ is a strict local maximum of F_g.*

Theorem 3 (Necessity)**.** *Let $g : [0, \infty) \to \mathbb{R}$ be a continuous function which is twice continuously differentiable on $(0, \infty)$, and let $F_g : S^{m-1} \to \mathbb{R}$ be constructed from g according to equation (1).*

1. *If P1 does not hold for g, then there exists an integer $m > 1$ and strictly positive values of the parameters α_i, β_i such that F_g has a local maximum not contained in the set $\{\pm Z_i\}_{i=1}^m$.*
2. *If P1 holds but P2 does not hold for g, there exist strictly positive values of the parameters α_i, β_i such that at least one of the canonical directions Z_i is not a local maximum for F_g.*

Spectral Clustering as Basis Recovery. It turns out that geometric basis recovery has direct implications for spectral clustering. In particular, when an n-vertex similarity graph has m connected components, the spectral embedding into $\mathbb{R}^m$ maps each vertex in the j^{th} connected component to a single point $y_j = \beta_j Z_j$ where $\beta_j = \|y_j\|$ and $Z_j = y_j / \|y_j\|$. It happens that the points $Z_1, \ldots, Z_m$ are orthogonal. Thus, letting x_i denote the embedded points and defining

$$F_g(u) := \frac{1}{n} \sum_{i=1}^n g(|\langle u, x_i \rangle|) \ ,$$

there exist strictly positive weights $\alpha_1, \ldots, \alpha_m$ such that $F_g(u) = \sum_{j=1}^m \alpha_j g(\beta_j |\langle u, Z_j \rangle|)$. In particular, α_j is the fraction of vertices contained in the j^{th} component. Recovery of the basis directions $\{\pm Z_j\}_{j=1}^m$ corresponds to the recovery of the component clusters.

As the weights α_j and β_j take on a special form in spectral clustering, it happens that property P1 by itself is sufficient to guarantee that the local maxima of F_g are precisely the basis directions $\{\pm Z_j\}_{j=1}^m$. As this article primarily focuses on the problem of spectral clustering, the main text will focus on the basis recovery problem arising in spectral clustering.

3 Spectral clustering: the Problem Statement

Let $G = (V, A)$ denote a similarity graph where V is a set of n vertices and A is an adjacency matrix with non-negative weights. Two vertices $i, j \in V$ are incident if $a_{ij} > 0$, and the value of a_{ij} is interpreted as a measure of the similarity between the vertices. In spectral clustering, the goal is to partition the vertices of a graph into sets $\mathcal{S}_1, \ldots, \mathcal{S}_m$ such that these sets form natural clusters in the graph. In the most basic setting, G consists of m connected components, and the natural clusters should be the components themselves. In this case, if $i' \in \mathcal{S}_i$ and $j' \in \mathcal{S}_j$ then $a_{i'j'} = 0$ whenever $i \neq j$. For convenience, we can consider the vertices of V to be indexed such that all indices in $\mathcal{S}_i$ precede all indices in $\mathcal{S}_j$ when $i < j$. Matrix A takes on the form:

$$A = \begin{pmatrix} A_{\mathcal{S}_1} & 0 & \cdots & 0 \\ 0 & A_{\mathcal{S}_2} & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & A_{\mathcal{S}_m} \end{pmatrix},$$

a block diagonal matrix. In this setting, spectral clustering can be viewed as a technique for reorganizing a given similarity matrix A into such a block diagonal matrix.

In practice, G rarely consists of m truly disjoint connected components. Instead, one typically observes a matrix $\tilde{A} = A + E$ where E is some error matrix with (hopefully small) entries e_{ij}. For i and j in different clusters, all that can be said is that $\tilde{a}_{ij}$ should be small. The goal of spectral clustering is to permute the rows and columns of $\tilde{A}$ to form a matrix which is nearly block diagonal and to recover the corresponding clusters.

4 Graph Laplacian's Null Space Structure

Given an n-vertex similarity graph $G = (V, A)$, define the diagonal degree matrix D with non-zero entries $d_{ii} = \sum_{j \in V} a_{ij}$. The unnormalized Graph Laplacian is defined as $L := D - A$. The following well known property of the graph Laplacian (see [21] for a review) helps shed light on its importance: Given $u \in \mathbb{R}^n$,

$$u^T L u = \frac{1}{2} \sum_{i,j \in V} a_{ij}(u_i - u_j)^2 \ . \tag{2}$$

The graph Laplacian L is symmetric positive semi-definite as equation (2) cannot be negative. Further, u is a 0-eigenvector of L (or equivalently, $u \in \mathcal{N}(L)$) if and only if $u^T L u = 0$. When G consists of m connected components with indices in the sets $\mathcal{S}_1, \ldots, \mathcal{S}_m$, inspection of equation (2) gives that $u \in \mathcal{N}(L)$ precisely when u is piecewise constant on each $\mathcal{S}_i$. In particular,

$$\{|\mathcal{S}_1|^{-1/2}\mathbf{1}_{\mathcal{S}_1}, \ldots, |\mathcal{S}_m|^{-1/2}\mathbf{1}_{\mathcal{S}_m}\}$$

is an orthonormal basis for $\mathcal{N}(L)$.

In general, letting $X \in \mathbb{R}^{d \times m}$ contain an orthogonal basis of $\mathcal{N}(L)$, it cannot be guaranteed that the rows of X will act as indicator vectors for the classes, as the columns of X can only been characterized up to a rotation within the subspace $\mathcal{N}(L)$. However, the rows of X are contained in a scaled orthogonal basis of $\mathbb{R}^m$ with the basis directions corresponding to the various classes. We use the following formulation of this result (see [22], [20, Proposition 5], and [16, Proposition 1] for related statements).

Proposition 4. *Let the similarity graph $G = (V, A)$ contain m connected components with indices in the sets $\mathcal{S}_1, \ldots, \mathcal{S}_m$, let $n = |V|$, and let L be the unnormalized graph Laplacian of G. Then, $\mathcal{N}(L)$ has dimensionality m. Let $X = (x_{\cdot 1}, \ldots, x_{\cdot m})$ contain m scaled, orthogonal column vectors forming*

a basis of $\mathcal{N}(L)$ *such that* $\|x_{\cdot j}\| = \sqrt{n}$ *for each* $j \in [m]$*. Then, there exist weights* $w_1, \ldots, w_m$ *with* $w_j = \frac{|\mathcal{S}_j|}{n}$ *and mutually orthogonal vectors* $Z_1, \ldots, Z_m \in \mathbb{R}^m$ *such that whenever* $i \in \mathcal{S}_j$*, the row vector* $x_{i\cdot} = \frac{1}{\sqrt{w_j}} Z_j^T$*.*

Proof. We define the matrix $M_{\mathcal{S}_i} := \mathbf{1}_{\mathcal{S}_i} \mathbf{1}_{\mathcal{S}_i}^T$. $P_{\mathcal{N}(L)}$ can be constructed from any orthonormal basis of $\mathcal{N}(L)$. In particular, using the two bases $\{|\mathcal{S}_1|^{-1/2} \mathbf{1}_{\mathcal{S}_1}, \ldots, |\mathcal{S}_m|^{-1/2} \mathbf{1}_{\mathcal{S}_m}\}$ and $\{\frac{1}{\sqrt{n}} x_{\cdot 1}, \ldots, \frac{1}{\sqrt{n}} x_{\cdot m}\}$ yields:

$$P_{\mathcal{N}(L)} = \sum_{i=1}^{m} |\mathcal{S}_i|^{-1} M_{\mathcal{S}_i} \quad \text{and} \quad P_{\mathcal{N}(L)} = \frac{1}{n} X X^T \ .$$

Thus for $i, j \in V$, $\frac{1}{n} \langle x_{i\cdot}, x_{j\cdot} \rangle = (P_{\mathcal{N}(L)})_{ij}$. In particular, if there exists $\ell \in [m]$ such that $i, j \in \mathcal{S}_\ell$, then $\frac{1}{n} \langle x_{i\cdot}, x_{j\cdot} \rangle = |\mathcal{S}_\ell|^{-1}$. When i and j belong to separate clusters, then $x_{i\cdot} \perp x_{j\cdot}$.

For i, j in component ℓ,

$$\cos(\angle(x_{i\cdot}, x_{j\cdot})) = \frac{\langle x_{i\cdot}, x_{j\cdot} \rangle}{\|x_{i\cdot}\| \|x_{j\cdot}\|} = \frac{|S_\ell|^{-1}}{|S_\ell|^{-1/2} |S_\ell|^{-1/2}} = 1 \ ,$$

giving that $x_{i\cdot}$ and $x_{j\cdot}$ are in the same direction. As they have the same magnitude as well, $x_{i\cdot}$ and $x_{j\cdot}$ coincide for any two indices i and j belonging to the same component of G.

Thus letting $w_i = \frac{|\mathcal{S}_i|}{n}$ for $i = 1, \ldots, m$, there are m perpendicular vectors $Z_1, \ldots, Z_m$ corresponding to the m connected components of G such that $x_{i\cdot} = \frac{1}{\sqrt{w_\ell}} Z_\ell^T$ for each $i \in \mathcal{S}_\ell$. □

Proposition 4 demonstrates that using the null space of the unnormalized graph Laplacian, the m connected components in G are mapped to m basis vectors in $\mathbb{R}^m$. Of course, under a perturbation of A, the interpretation of Proposition 4 must change. In particular, G will no longer consist of m connected components, and instead of using only vectors in $\mathcal{N}(L)$, X must be constructed using the eigenvectors corresponding to the lowest m eigenvalues of L. With the perturbation of A comes a corresponding perturbation of these eigenvectors. When the perturbation is not too large, the resulting rows of X yield m nearly orthogonal clouds of points.

Due to different properties of the resulting spectral embeddings, normalized graph Laplacians are often used in place of L for spectral clustering, in particular $L_{\text{sym}} := D^{-1/2} L D^{-1/2}$ and $L_{\text{rw}} := D^{-1} L$. Using either L_{sym} or L_{rw} in place of L when generating the eigenvector matrix X from Proposition 4 will be fully consistent with the proposed algorithms of this paper. When G consists of m connected components, $\mathcal{N}(L_{\text{rw}})$ happens to be the same as $\mathcal{N}(L)$, making Proposition 4 and all subsequent results in this paper equally applicable to L_{rw}. If L_{sym} is used, the scaled basis structure created in Proposition 4 is replaced by a slightly more involved ray structure. The admissibility of L_{sym} is discussed in Appendix B.

In the perturbed setting, it is natural to define a notion of a best clustering within a similarity graph. This can be done using graph cuts. For two sets of vertices $\mathcal{S}_1$ and $\mathcal{S}_2$, the cut is defined as $\text{Cut}(\mathcal{S}_1, \mathcal{S}_2) := \sum_{i \in \mathcal{S}_1,\ j \in \mathcal{S}_2} a_{ij}$. In the m-way min cut problem, the goal is to partition the vertices into m non-empty sets which minimize the cost

$$C_{\text{Cut}}(\{\mathcal{S}_1, \ldots, \mathcal{S}_m\}) := \sum_{i=1}^{m} \text{Cut}(\mathcal{S}_i, \mathcal{S}_i^c) \ .$$

Such a partition gives an "optimal" set of m clusters. However, this "optimal" set of clusters does not penalize small clusters, making it plausible that one would find an "optimal" cut which simply detaches $m - 1$ nearly isolated vertices or small clusters from the rest of the graph.

In order to favor clusters of more equal size, one can consider variants of the m-way min cut problem. In particular, spectral clustering arises as the relaxation to the following NP-hard graph cut problems: the m-way min normalized cut and the m-way min ratio cut. In the ratio cut problem, one minimizes the cost

$$C_{\mathrm{RCut}}(\{\mathcal{S}_1, \cdots, \mathcal{S}_m\}) := \sum_{i=1}^{m} \frac{\mathrm{Cut}(\mathcal{S}_i, \mathcal{S}_i^c)}{|\mathcal{S}_i|} \ .$$

Alternatively, in the normalized m-way cut problem, one minimizes the cost

$$C_{\mathrm{NCut}}(\{\mathcal{S}_1, \cdots, \mathcal{S}_m\}) := \sum_{i=1}^{m} \frac{\mathrm{Cut}(\mathcal{S}_i, \mathcal{S}_i^c)}{\sum_{j \in \mathcal{S}_i} d_{jj}} \ .$$

Spectral clustering using L arises as a relaxation of the m-way ratio cut problem (see [21] section 5), whereas the using of L_{sym} arises as a relaxation of the m-way normalized cut problem (see [26]). These interpretations do not give error bounds under perturbation, but do give some insight as to how clusters are formed. For instance, using L_{sym} discourages classifying a few sparsely connected points as a cluster since the normalization term $\sum_{j \in \mathcal{S}_i} d_{jj}$ for that cluster in C_{NCut} is also small.

On the other hand, the use of L_{rw} arises from the theory of random walks. The matrix $D^{-1}A$ has rows of unit norm, allowing the entry $(D^{-1}A)_{ij}$ to be interpreted as the transition probability from state i to state j in a Markov random chain. As $L_{\mathrm{rw}} = I - D^{-1}A$, the smallest eigenvectors of L_{rw} correspond to the largest eigenvectors of $D^{-1}A$, making eigenvector analysis of L_{rw} formally equivalent to eigenvector analysis of the stochastic matrix $D^{-1}A$. In this setting, the notion of m clusters is recast as having a random walk with m nearly invariant aggregates (see e.g. references [6, 15]). In particular, when G consists of m connected components, there are m sets or aggregates $\mathcal{S}_1, \ldots, \mathcal{S}_m$ of states that are truly invariant, meaning that the transition probabilities between states belonging to distinct sets is 0. Creating X as in Proposition 4 using L_{rw} in place of L maps each state within an aggregate to a single basis vector.

5 Basis Recovery for Spectral Clustering

Interpreting the Embedding's Basis Structure. Given a graph G with n vertices and m connected components, let X; $\mathcal{S}_1, \ldots, \mathcal{S}_m$; $w_1, \ldots, w_m$; and $Z_1, \ldots, Z_m$ be constructed from L as in Proposition 4. The basis vectors $Z_1, \ldots, Z_m$ are mutually orthogonal in $\mathbb{R}^m$, and each weight $w_i = \frac{|\mathcal{S}_i|}{n}$ is the fraction of the rows of X indexed as $x_{\ell \cdot}$ coinciding with the point $\frac{1}{\sqrt{w_i}} Z_i^T$. It suffices to recover the basis directions Z_i up to sign in order to cluster the points. That is, each embedded point $x_{j \cdot} \in \mathcal{S}_i$ lies on the line through $\pm Z_i$ and the origin, making these lines correspond to the clusters.

We use an approach based on function optimization over projections of the embedded data. Let $F_g : S^{m-1} \to \mathbb{R}$ be defined on the unit sphere in terms of a "contrast function" $g : [0, \infty) \to \mathbb{R}$ as follows:

$$F_g(u) := \frac{1}{n} \sum_{i=1}^{n} g(|\langle u, x_{i \cdot} \rangle|) \tag{3}$$

This can equivalently be written:

$$F_g(u) = \sum_{i=1}^{m} w_i g(\tfrac{1}{\sqrt{w_i}} |\langle u, Z_i \rangle|) \ . \tag{4}$$

In equation (4), F_g takes on a special form of the basis recovery problem presented in equation (1) with the choices $\alpha_i = w_i$ and $\beta_i = \frac{1}{\sqrt{w_i}}$. Due to the special form of these weights, only property P1 is required in order to recover the directions $\{\pm Z_i : i \in [m]\}$:

Theorem 5. *Let $g : [0, \infty) \to \mathbb{R}$ be a continuous function satisfying property P1. Let $F_g : S^{m-1} \to \mathbb{R}$ be defined from g according to equation (4). Then, the set $\{\pm Z_i : i \in [m]\}$ is a complete enumeration of the local maxima of F_g.*

Stability analysis: It can be shown that both the embedding structure (Proposition 4) and the local maxima structure of F_g (Theorem 5) are robust to a perturbation from the setting in which G consists of m connected components. We provide such a stability analysis, demonstrating that our algorithms are robust to such perturbations. However, since these results are very technical in nature, they are deferred to the appendices D and E in order to simplify the exposition.

Theorem 5 implies that a function optimization problem defined using the spectral embedding of L can be used to recover the clusters corresponding to the connected components of G. It should be noted that L_{sym} can similarly be used for spectral clustering. Suppose that L_{sym} is constructed from G. Then, $\dim(\mathcal{N}(L_{\text{sym}})) = m$. Further, it can be shown that if X' contains a scaled, orthogonal basis of $\mathcal{N}(L_{\text{sym}})$ in its columns such that $\|x'_{\cdot i}\| = \sqrt{n}$, then there exists an orthonormal basis $Z'_1, \ldots, Z'_m$ of $\mathbb{R}^m$ such that $\angle(Z'_i, x'_{j\cdot}) = 0$ for each $j \in \mathcal{S}_i$. Defining $F_g^{\text{sym}} : S^{m-1} \to \mathbb{R}$ from g by

$$F_g^{\text{sym}}(u) := \frac{1}{n} \sum_{i=1}^{n} g(|\langle u, x'_{i\cdot} \rangle|) \tag{5}$$

then we have the following result for spectral clustering using L_{sym}:

Theorem 6. *Let $g : [0, \infty) \to \mathbb{R}$ be a continuous function satisfying property P1. Let $F_g^{\text{sym}} : S^{m-1} \to \mathbb{R}$ be defined from g according to equation (5). Then, $\{\pm Z'_i : i \in [m]\}$ is a complete enumeration of the local maxima of F_g^{sym}.*

We focus on spectral clustering using L in order to simplify the exposition. Nevertheless, the proof of Theorem 6 is provided in Appendix B.

We now proceed to prove Theorem 5. As $Z_1, \ldots, Z_m$ form an orthonormal basis of the space, we will for simplicity work in the unknown coordinate system where $Z_1, \ldots, Z_m$ are the canonical vectors $e_1, \ldots, e_m$. The proof proceeds by establishing sufficient and necessary optimality conditions in a series of Lemmas exploiting the convexity structure induced by the substitution $u_i \mapsto u_i^2$ which maps the domain S^{m-1} onto the simplex $\Delta^{m-1} := \text{conv}(e_1, \ldots, e_m)$.

Lemma 7. *Let $h_i : [0, \infty) \to \mathbb{R}$, $i \in [m]$ be a family of strictly convex functions. Let $H : [0, \infty)^m \to \mathbb{R}$ be given by $H(t) = \sum_{i=1}^m h_i(t_i)$. Let $c_i > 0$ for $i \in [m]$. Then the set of local maxima of H relative to the scaled simplex $\text{conv}\{c_i e_i\}_{i=1}^m$ is contained in the set $\{c_i e_i\}_{i=1}^m$.*

Proof. We have that H is strictly convex in its domain and $\{w_i e_i\}_{i=1}^m$ is the set of extreme points of $\text{conv}\{c_i e_i\}_{i=1}^m$. The result follows immediately from Fact 1. □

Recalling that $Z_1, \ldots, Z_m$ has been identified with the canonical vectors of our unknown coordinate system, consider equation (4). By defining $g_i : [0, \infty) \to \mathbb{R}$ by $g_i(t) := w_i g((1/\sqrt{w_i})u_i)$, then we obtain $F_g(u) = \sum_{i=1}^m g_i(|u_i|)$ and $t \mapsto g_i(\sqrt{t})$ is a strictly convex function for each $i \in [m]$. Thus, the following three Lemmas largely demonstrate why Theorem 5 holds.

Lemma 8 (Necessary optimality condition)**.** *Let $g_i : [0, \infty) \to \mathbb{R}$, $i = 1, \ldots, m$ be such that $t \mapsto g_i(\sqrt{t})$ is strictly convex in $[0, \infty)$. Let $F_g : \mathbb{R}^m \to \mathbb{R}$ be given by $F_g(u) = \sum_i g_i(|u_i|)$. Then the set of local maxima of F_g relative to S^{m-1} is contained in $\{\pm e_i : i \in [m]\}$.*

Proof. Suppose $z \in S^{m-1}$ is not in $\{\pm e_i : i \in [m]\}$. Without loss of generality we assume $z_i \geq 0$ for all $i \in [m]$. Let $h_i : [0,\infty) \to \mathbb{R}$, $i \in [m]$ be given by $h_i(t) = g_i(\sqrt{t})$. Let $H : [0,\infty)^m \to \mathbb{R}$ be given by $H(y) = \sum_{i=1}^m h_i(y_i)$, as in Lemma 7. By definition we have that H is a sum of convex functions and hence is convex. Moreover, it is easy to see that it is strictly convex in its domain. The map $G : [0,\infty)^m \to [0,\infty)^m$ given by $u \mapsto \psi(u) = (u_i^2)$ is a homeomorphism between $S^{m-1} \cap [0,\infty)^m$ and $\Delta^{m-1} := \text{conv}\{e_i : i \in [m]\}$. We also have $F_g(u) = H(\psi(u))$ in S^{m-1}. This implies that the properties of being a local maximum of H relative to Δ^{m-1} and being a local maximum of F_g relative to $S^{m-1} \cap [0,\infty)^m$ are preserved under ψ, and local maxima are in one-to-one correspondence. As $\psi(z)$ is not a canonical vector, Lemma 7 implies that $\psi(z)$ is not a local maximum of H relative to Δ^{m-1}. Thus, z is not a local maximum of F_g relative to $S^{m-1} \cap [0,\infty)^m$, in particular, also not relative to S^{m-1}. □

Lemma 9. *Let $h : [0,\infty) \to \mathbb{R}$ be a strictly convex function. Let $w_i > 0$ for $i \in [m]$. Let $H : [0,\infty)^m \to \mathbb{R}$ be given by $H(x) = \sum_i w_i h(x_i/w_i)$. Then the set $\{e_i\}_{i=1}^m$ is contained in the set of strict local maxima of H relative to Δ^{m-1}.*

(Note that this lemma is written with greater generality: We do not assume that $\sum w_i = 1$.)

Proof. By symmetry, it is enough to show that e_1 is a strict local maximum. Let $h_i(t) = w_i h(t/w_i)$. In this way we have $H(x) = \sum_{i=1}^m h_i(x_i)$. We need to understand the behavior of h_1 around 1 and $h_2, \ldots, h_m$ around 0. To this end and to take advantage of strict convexity, we will consider a (two piece) piecewise affine interpolating upper bound to each h_i. We will pick $t_i \in (0,1)$ so that the approximation is affine in $[0,t_i]$ and $[t_i,1]$ with t_i so that the slope of the left piece is the same for all i. The slope of the right piece is always larger than the slope of the left piece by strict convexity. We actually choose t_i to be at most $1/2m$ so that the behavior of H in a neighborhood of e_1 in Δ^{m-1} is controlled by the right piece $[t_i,1]$ for h_1 and by the left piece $[0,t_i]$ for $h_2, \ldots, h_m$. We make these choices more precise now. Let $w_{\max} = \max\{w_i : i \in [m]\}$. Let $t_i = w_i/(2m w_{\max})$. We have $0 < t_i \leq 1/(2m)$. The piecewise affine upper bound to h_i is the interpolant through $0, t_i, 1$. The left piece has slope

$$\frac{h_i(t_i) - h_i(0)}{t_i} = \frac{w_i h(\frac{1}{2m w_{\max}}) - w_i h(0)}{\frac{w_i}{2m w_{\max}}} = \frac{h(\frac{1}{2m w_{\max}}) - h(0)}{\frac{1}{2m w_{\max}}},$$

which is independent of i and we denote m_l. The right piece of the interpolant of h_i has a slope that we denote m_i. By strict convexity we have $m_i > m_l$. The fact that the interpolating pieces are upper bounds implies the following inequalities:

$$\begin{aligned} h_i(x) &< m_l x + w_i h(0) \qquad \text{for } x \in (0,t_i), \\ h_i(x) &< h_i(1) - m_i(1-x) \quad \text{for } x \in (t_i,1). \end{aligned} \tag{6}$$

Consider the neighborhood N of e_1 relative to Δ^{m-1} given by $N = \{y \in \Delta^{m-1} : y_i \leq t_i \text{ for } i = 2, \ldots, m\}$. For $y \in N$ we have $y_1 \geq 1/2$. Putting everything together, for $y \in N \setminus \{e_1\}$ we have:

$$\begin{aligned} H(e_1) - H(y) &= h_1(1) + \sum_{i=2}^m w_i h(0) - \sum_{i=1}^m h_i(y_i) \\ &\geq h_1(1) + \sum_{i=2}^m w_i h(0) - [h_1(1) - m_1(1-y_1) + \sum_{i=2}^m (m_l y_i + w_i h(0))] \\ &\qquad \text{(using (6))} \\ &= m_1(1-y_1) - m_l \sum_{i=2}^m y_i = (m_1 - m_l)(1-y_1) > 0. \end{aligned}$$

□

Lemma 10 (Sufficient optimality condition)**.** *Let $g : [0, \infty) \to \mathbb{R}$ be such that $t \mapsto g(\sqrt{t})$ is strictly convex in $[0, \infty)$. Let $w_i > 0$ for $i \in [m]$. Let $F_g : \mathbb{R}^m \to \mathbb{R}$ be given by $F_g(u) = \sum_i w_i g(|u_i|/\sqrt{w_i})$. Then the set $\{\pm e_i : i \in [m]\}$ is contained in the set of strict local maxima of G relative to S^{m-1}.*

Proof. By symmetry, it is enough to show that e_1 is a strict local maximum of F_g relative to $S^{m-1} \cap [0, \infty)^m$. The homeomorphism from the proof of Lemma 8 implies that it is enough to show that e_1 is a strict local maximum of $H(x) = \sum_i w_i h(x_i/w_i)$ relative to Δ^{m-1}. This follows immediately from Lemma 9. □

The proof of the main result (Theorem 5) now follows quite easily:

Proof of Theorem 5. That $B = \{\pm m_i : i \in [m]\}$ give strict local maxima of F_g follows from Lemma 10. To see that F_g has no other local maxima besides those in B, use Lemma 8. □

6 Spectral Clustering Algorithms

6.1 Choosing a contrast function

There are many possible choices of contrast g which are admissible for spectral clustering under Theorems 5 and 6 including the following:

$$g_p(t) = |t|^p \text{ where } p \in (2, \infty) \qquad g_{abs}(t) = -|t| \qquad g_{ht}(t) = \log\cosh(t)$$
$$g_{sig}(t) = -\frac{1}{1 + \exp(-|t|)} \qquad g_{gau} = e^{-t^2} \qquad .$$

In choosing contrasts, it is instructive to first consider the function $g_2(y) = y^2$ (which fails to satisfy property P1 and is thus not admissible). Noting that $F_{g_2}(u) = \sum_{i=1}^m w_i(\frac{1}{\sqrt{w_i}}\langle u, Z_i\rangle)^2 = 1$, we see that F_{g_2} is constant on the unit sphere. We see that the distinguishing power of a contrast function for spectral clustering comes from property P1. Intuitively, "more convex" contrasts g have better resolving power but are also more sensitive to outliers and perturbations of the data. Indeed, if g grows rapidly, a small number of outliers far from the origin could significantly distort the maxima structure of F_g.

Due to this tradeoff, g_{sig} and g_{abs} could be important practical choices for the contrast function. Both $g_{sig}(\sqrt{x})$ and $g_{abs}(\sqrt{x})$ have a strong convexity structure near the origin. As g_{sig} is a bounded function, it should be very robust to perturbations. In comparison, $g_{abs}(\sqrt{t}) = -|\sqrt{t}|$ maintains a stronger convexity structure over a much larger region of its domain and has only a linear rate of growth as $n \to \infty$. This is a much slower growth rate than is present for instances in g_p with $p > 2$.

6.2 Algorithms

We now have all the tools needed to create a new class of algorithms for spectral clustering. Given a similarity graph $G = (V, A)$ containing n vertices, define a graph Laplacian $\tilde{L}$ among L, L_{rw}, and L_{sym} (reader's choice). Viewing G as a perturbation of a graph consisting of m connected components, construct $X \in \mathbb{R}^{n \times m}$ such that $x_{\cdot i}$ gives the eigenvector corresponding to the i^{th} smallest eigenvalue of $\tilde{L}$ with scaling $\|x_{\cdot i}\| = \sqrt{n}$.

With X in hand, choose a contrast function g satisfying P1. From g, the function $F_g(u) = \frac{1}{n}\sum_{i=1}^n g(\langle u, x_{i\cdot}\rangle)$ is defined on S^{m-1} using the rows of X. The local maxima of F_g correspond to the desired clusters of the graph vertices. Since F_g is a symmetric function, if F_g has a local maximum at u, F_g also has a local maximum at $-u$. However, the directions u and $-u$ correspond to the same line through the origin in $\mathbb{R}^m$ and form an equivalence class, with each such equivalence class corresponding to a cluster.

Our first goal is to find local maxima of F_g corresponding to distinct equivalence classes. We will use that the desired maxima of F_g should be approximately orthogonal to each other. Once we have obtained local maxima $u_1, \ldots, u_m$ of F_g, we cluster the vertices of G by placing vertex i in the j^{th} cluster using the rule $j = \arg\max_\ell |\langle u_\ell, x_{i\cdot}\rangle|$. We sketch two algorithmic ideas in HBROPT and HBRENUM. Here, HBR stands for hidden basis recovery.

Algorithm 1 Finds the local maxima of F_g defined from the points $x_{i\cdot}$ needed for clustering. The second input η is the learning rate (step size).

1: **function** HBROPT(X, η)
2: $C \leftarrow \{\}$
3: **for** $i \leftarrow 1$ to m **do**
4: Draw u from $S^{m-1} \cap \text{span}(C)^\perp$ uniformly at random.
5: **repeat**
6: $u \leftarrow u + \eta(\nabla F_g(u) - uu^T \nabla F_g(u)) \quad (= u + \eta P_{u^\perp} \nabla F_g(u))$
7: $u \leftarrow P_{\text{span}(C)^\perp} u$
8: $u \leftarrow \frac{u}{\|u\|}$
9: **until** Convergence
10: Let $C \leftarrow C \cup \{u\}$
11: **return** C

HBROPT is a form of projected gradient ascent. The parameter η is the learning rate. Each iteration of the repeat-until loop moves u in the direction of steepest ascent. For gradient ascent in $\mathbb{R}^m$, one would expect step 6 of HBROPT to read $u \leftarrow u + \eta \nabla F_g(u)$. However, gradient ascent is being performed for a function F_g defined on the unit sphere, but the gradient described by ∇F_g is for the function F_g with domain $\mathbb{R}^m$. The more expanded formula $\nabla F_g(u) - uu^T \nabla F_g(u)$ is the projection of ∇F_g onto the tangent plane of S^{m-1} at u. This update keeps u near the sphere.

We may draw u uniformly at random from $S^{m-1} \cap \text{span}(C)^\perp$ by first drawing u from S^{m-1} uniformly at random, projecting u onto $\text{span}(C)^\perp$, and then normalizing u. It is important that u stay near the orthogonal complement of $\text{span}(C)$ in order to converge to a new cluster rather than converging to a previously found optimum of F_g. Step 7 enforces this constraint during the update step.

Algorithm 2 Finds the local maxima of F_g defined from the points $x_{i\cdot}$ needed for clustering. The second input δ controls how far a point needs to be from previously found cluster centers to be a candidate future cluster center.

1: **function** HBRENUM(X, δ)
2: $C \leftarrow \{\}$
3: **while** $|C| < m$ **do**
4: $j \leftarrow \arg\max_i \{F_g(\frac{x_{i\cdot}}{\|x_{i\cdot}\|}) : \angle(\frac{x_{i\cdot}}{\|x_{i\cdot}\|}, u) > \delta \ \forall u \in C\}$
5: $C \leftarrow C \cup \{\frac{x_{j\cdot}}{\|x_{j\cdot}\|}\}$
6: **return** C

In contrast to HBROPT, HBRENUM more directly uses the point separation implied by the orthogonality of the approximate cluster centers. Since each embedded data point should be near to a cluster center, the data points themselves are used as test points. Instead of directly enforcing orthogonality between cluster means, a parameter $\delta > 0$ specifies the minimum allowable angle between found cluster means.

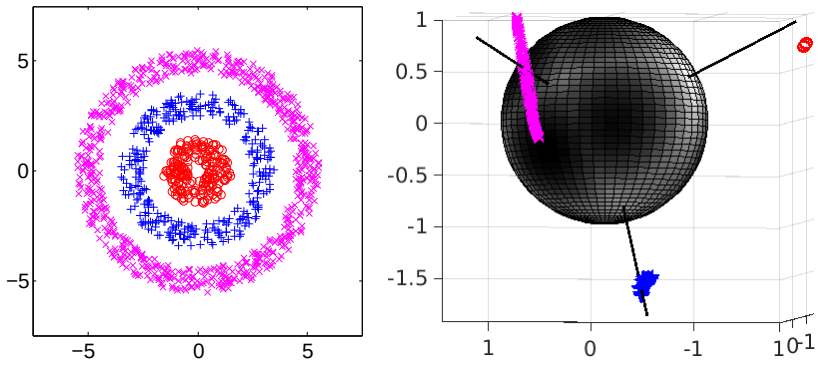

(a) (b)

Figure 1: An illustration of spectral clustering on the concentric circle data. (a) The output of clustering. (b) The embedded data and the contrast function.

By pre-computing the values of $F_g(x_{i\cdot}/\|x_{i\cdot}\|)$ outside of the while loop, HBRENUM can be run in $O(mn^2)$ time. For large similarity graphs, HBRENUM is likely to be slower than HBROPT which takes $O(m^2nt)$ time where t is the average number of iterations to convergence. The number of clusters m cannot exceed (and is usually much smaller than) the number of graph vertices n.

HBRENUM has a couple of nice features which may make it preferable on smaller data sets. Each center found by HBRENUM will always be within a cluster of data points even when the optimization landscape is distorted under perturbation. In addition, the maxima found by HBRENUM are based on a more global outlook, which may be important in the noisy setting.

7 Experiments

7.1 An Illustrating Example

Figure 1 illustrates our function optimization framework applied to spectral clustering. In this example, random points $p_1, p_2, \ldots, p_{1250}$ were generated from 3 concentric circles: 200 points were drawn uniformly at random from a radius 1 circle, 350 points from a radius 3 circle, and 700 points from a radius 5 circle. The points were then radially perturbed. The generated points are displayed in Figure 1a. From this data, a similarity matrix A was constructed as $a_{ij} = \exp(-\frac{1}{4}\|p_i - p_j\|^2)$, and the Laplacian embedding was performed using L_{rw}.

Figure 1b depicts the clustering process with the contrast g_{sig} on the resulting embedded points. In this depiction, the embedded data sufficiently encodes the desired basis structure that all local maxima of $F_{g_{sig}}$ correspond to desired clusters. The value of $F_{g_{sig}}$ is displayed by the grayscale heat map on the unit sphere in Figure 1b, with lighter shades of gray indicate greater values of $F_{g_{sig}}$. The cluster labels were produced using HBROPT. The rays protruding from the sphere correspond to the basis directions recovered by HBROPT, and the recovered labels are indicated by the color and symbol used to display each data point.

7.2 Image Segmentation Examples

Spectral clustering was first applied to image segmentation in [19], and it has remained a popular application of spectral clustering. The goal in image segmentation is to divide an image into regions which represent distinct objects or features of the image. Figure 2 illustrates segmentations produced by HBROPT-g_{abs} and spherical k-means on several example images from the BSDS300 test set [14]. For these images, the similarity matrix was constructed using only color and proximity

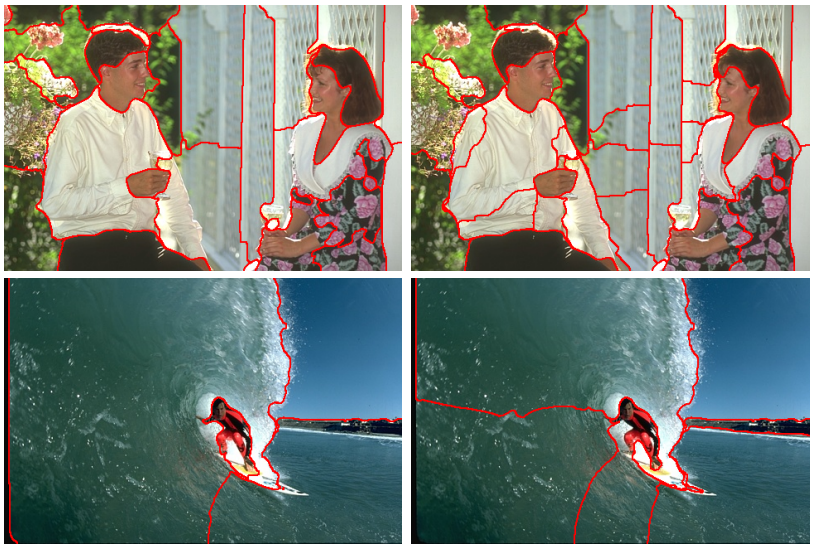

Figure 2: Segmented images from the BSDS300 test set. Red pixels mark the borders between segmented regions. Segmentation using HBRopt-g_{abs} (left panels) compared to spherical k-means (right panels).

	oracle-centroids	k-means-cosine	HBRopt					HBRenum				
			g_{abs}	g_{gau}	g_3	g_{ht}	g_{sig}	g_{abs}	g_{gau}	g_3	g_{ht}	g_{sig}
E. coli	79.7	69.0	80.9	81.2	79.3	81.2	80.6	68.7	**81.5**	**81.5**	68.7	**81.5**
Flags	33.2	33.1	**36.8**	34.1	36.6	**36.8**	34.4	34.7	**36.8**	**36.8**	34.7	**36.8**
Glass	49.3	46.8	**47.0**	46.8	**47.0**	**47.0**	46.8	**47.0**	**47.0**	**47.0**	**47.0**	**47.0**
Thyroid Disease	72.4	80.4	**82.4**	81.3	82.2	82.2	81.5	81.8	82.2	82.2	81.8	82.2
Car Evaluation	56.1	36.4	37.0	36.3	36.3	35.2	36.6	49.6	32.3	41.1	**49.9**	41.1
Cell Cycle	74.2	62.7	64.3	64.4	63.8	64.5	64.0	60.1	62.9	**64.8**	61.1	62.7

Table 1: Percentage accuracy of spectral clustering algorithms. The best performing contrast function among the HBRopt and HBRenum algorithms for each data set is bolded.

information.

7.3 Stochastic block model with imbalanced clusters

In this example, we construct a similarity graph $A = \text{diag}(A_1, A_2, A_3) + E$ where each A_i is a symmetric matrix corresponding to a cluster and E is a small perturbation. We set $A_1 = A_2$ to be 10×10 matrices with entries 0.1. We set A_3 to be a 1000×1000 matrix which is symmetric, approximately 95% sparse with randomly chosen non-zero locations set to 0.001. When performing this experiment 50 times, HBRopt-g_{sig} obtained a mean accuracy of 99.9%. This is in contrast to spherical k-means with randomly chosen starting points which obtained a mean accuracy of only 42.1%. This disparity is due to the fact that splitting the large cluster is in fact optimal in terms of the spherical k-means objective function but leads to poor classification performance. Our method does not suffer from that shortcoming.

7.4 Performance Evaluation on UCI datasets.

We compare spectral clustering performance on a number of data sets with unbalanced cluster sizes. In particular, the E. coli, Flags, Glass, Thyroid Disease, and Car Evaluation data sets which are part of the UCI machine learning repository [3] are used. We also use the standardized gene expression data set [24, 25], which is also referred to as Cell Cycle. For the Flags data set, we used religion as the ground truth labels, and for Thyroid Disease, we used the new-thyroid data.

For all data sets, we only used fields for which there were not missing values, we normalized the data such that every field had unit standard deviation, and we constructed the similarity matrix A using a Gaussian kernel $k(y_i, y_j) = \exp(-\alpha\|y_i - y_j\|^2)$. The parameter α was chosen separately for each data set in order to create a good embedding. The choices of α were: 0.25 for E. Coli, 32 for Glass, 32 for Thyroid Disease, 128 for Flags, 0.25 for Car Evaluation, and 0.125 for Cell Cycle.

The spectral embedding was performed using the symmetric normalized Laplacian L_{sym}. Then, the clustering performance of our proposed algorithms HBRopt and HBRenum (implemented with $\delta = 3\pi/8$ radians) were compared with the following baselines:

- oracle-centroids: The ground truth labels are used to set means $\mu_j = \frac{1}{|\mathcal{S}_j|}\sum_{i\in\mathcal{S}_j}\frac{x_{i\cdot}}{\|x_{i\cdot}\|}$ for each $j \in [m]$. Points are assigned to their nearest cluster mean in cosine distance.

- k-means-cosine: The spherical k-means algorithm is run with a random initialization of the means, cf. [16].

We report the clustering accuracy of each algorithm in Table 1. The accuracy is computed using the best matching between the clusters and the true labels. The reported results consist of the mean performance over a set of 25 runs for each algorithm. The number of clusters being searched for was set to the ground truth number of clusters. In most cases, we see improvement in performance over spherical k-means.

8 Acknowledgments

This work was supported by National Science Foundation grant IIS 1117707.

References

[1] J. Anderson, N. Goyal, and L. Rademacher. Efficient learning of simplices. In S. Shalev-Shwartz and I. Steinwart, editors, *COLT*, volume 30 of *JMLR Proceedings*, pages 1020–1045. JMLR.org, 2013.

[2] F. R. Bach and M. I. Jordan. Learning spectral clustering, with application to speech separation. *Journal of Machine Learning Research*, 7:1963–2001, 2006.

[3] K. Bache and M. Lichman. UCI machine learning repository, 2013.

[4] M. Belkin and P. Niyogi. Laplacian eigenmaps for dimensionality reduction and data representation. *Neural Comput.*, 15(6):1373–1396, 2003.

[5] C. Davis and W. M. Kahan. The rotation of eigenvectors by a perturbation. iii. *SIAM Journal on Numerical Analysis*, 7(1):1–46, 1970.

[6] P. Deuflhard, W. Huisinga, A. Fischer, and C. Schütte. Identification of almost invariant aggregates in reversible nearly uncoupled markov chains. *Linear Algebra and its Applications*, 315(1):39–59, 2000.

[7] P. Gritzmann and V. Klee. On the 0–1-maximization of positive definite quadratic forms. In *Operations Research Proceedings 1988*, pages 222–227. Springer, 1989.

[8] J.-B. Hiriart-Urruty and C. Lemaréchal. *Convex Analysis and Minimization Algorithms: Part 1: Fundamentals*, volume 1. Springer, 1996.

[9] R. A. Horn and C. R. Johnson. *Topics in matrix analysis*. 1991.

[10] D. Hsu and S. M. Kakade. Learning mixtures of spherical Gaussians: moment methods and spectral decompositions. In *Proceedings of the 4th conference on Innovations in Theoretical Computer Science*, pages 11–20. ACM, 2013.

[11] A. Hyvärinen, J. Karhunen, and E. Oja. *Independent component analysis*, volume 46. John Wiley & Sons, 2004.

[12] A. K. Jain and R. C. Dubes. *Algorithms for clustering data*. Prentice-Hall, Inc., Upper Saddle River, NJ, USA, 1988.

[13] P. Kumar, N. Narasimhan, and B. Ravindran. Spectral clustering as mapping to a simplex. *2013 ICML workshop on Spectral Learning*, 2013.

[14] D. R. Martin, C. Fowlkes, D. Tal, and J. Malik. A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics. In *ICCV*, pages 416–425, 2001.

[15] M. Meilă and J. Shi. A random walks view of spectral segmentation. In *AI and Statistics (AISTATS)*, 2001.

[16] A. Y. Ng, M. I. Jordan, and Y. Weiss. On spectral clustering: Analysis and an algorithm. *Advances in neural information processing systems*, 2:849–856, 2002.

[17] R. T. Rockafellar. *Convex analysis*. Princeton Landmarks in Mathematics. Princeton University Press, Princeton, NJ, 1997. Reprint of the 1970 original, Princeton Paperbacks.

[18] L. Rosasco, M. Belkin, and E. D. Vito. On learning with integral operators. *The Journal of Machine Learning Research*, 11:905–934, 2010.

[19] J. Shi and J. Malik. Normalized cuts and image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22(8):888–905, 2000.

[20] D. Verma and M. Meilă. A comparison of spectral clustering algorithms. Technical report, University of Washington CSE Department, Seattle, WA 98195-2350, 2003. doi=10.1.1.57.6424, Accessed online via CiteSeerx 5 Mar 2014.

[21] U. Von Luxburg. A tutorial on spectral clustering. *Statistics and computing*, 17(4):395–416, 2007.

[22] M. Weber, W. Rungsarityotin, and A. Schliep. *Perron cluster analysis and its connection to graph partitioning for noisy data*. Konrad-Zuse-Zentrum für Informationstechnik Berlin, 2004.

[23] T. Wei. A study of the fixed points and spurious solutions of the deflation-based fastica algorithm. *Neural Computing and Applications*, pages 1–12, 2015.

[24] K. Y. Yeung, C. Fraley, A. Murua, A. E. Raftery, and W. L. Ruzzo. Model-based clustering and data transformations for gene expression data. *Bioinformatics*, 17(10):977–987, 2001.

[25] K. Y. Yeung, C. Fraley, A. Murua, A. E. Raftery, and W. L. Ruzzo. Model-based clustering and data transformations for gene expression data. `http://faculty.washington.edu/kayee/model/`, 2001. Accessed: 20 Jan 2015.

[26] S. X. Yu and J. Shi. Multiclass spectral clustering. In *Computer Vision, 2003. Proceedings. Ninth IEEE International Conference on*, pages 313–319. IEEE, 2003.

A Learning a Weighted Basis

In this appendix, we show that Theorems 2 and 3 hold. The key idea is that there is a natural isomorphism between a quadrant of the unit sphere and the simplex $\text{conv}(Z_1, \ldots, Z_m)$. As $Z_1, \ldots, Z_m$ gives an orthonormal basis of the space, we will without loss of generality work in the unknown coordinate system in which $Z_1, \ldots, Z_m$ is the canonical basis $e_1, \ldots, e_m$.

We let Q_1 be the first quadrant of the unit sphere, i.e. $Q_1 := S^{m-1} \cap [0, \infty)^m$. Then, $\psi : Q_1 \to \Delta^{m-1}$ defined by $\psi(u)_i = u_i^2$ is a homeomorphism (introduced earlier in the proof of Lemma 8). We define $H : \Delta^{m-1} \to \mathbb{R}$ as

$$H(t) := F_g \circ \psi^{-1}(t) = \sum_{i=1}^{m} \alpha_i g\left(\beta_i \sqrt{t_i}\right) \ ,$$

where F_g is defined in accordance with equation (1). Letting $h : [0, \infty) \to \infty$ be given by $h(x) = g(\sqrt{x})$, then, $H(t) = \sum_{i=1}^{m} \alpha_i h(\beta_i^2 t_i)$. Note that properties P1 and P2 hold for g if and only if the following properties hold for h:

P1*. The function h is strictly convex.

P2*. The (right) derivative at the origin $h'(x) = 0$ or $h'(0) = -\infty$.

As F_g and H take on the same values, it is clear that H takes on a local maximum for some $t \in \Delta^{m-1}$ if and only if F_g takes on a local maximum relative to Q_1 at $\psi^{-1}(t)$. The optima of F_g relative to S^{m-1} are fully determined by the optima of F_g relative to Q_1 and the symmetries of F_g. In this sense, optimization over the unit sphere is equivalent optimization over the unknown convex body Δ^{m-1}.

A.1 Sufficiency Argument for Weighted Basis Recovery

Lemma 11. *Suppose that $h : [0, \infty) \to \mathbb{R}$ satisfies properties P1* and P2*. For strictly positive constants $\alpha_1, \ldots, \alpha_m$ and $\beta_1, \ldots, \beta_m$, let $H : \Delta^{m-1} \to \mathbb{R}$ be given by $H(t) = \sum_{i=1}^{m} \alpha_i h(\beta_i^2 t_i)$. Then $\{e_i\}_{i=1}^{m}$ are local maxima of H.*

Proof. By symmetry, it suffices to show that e_1 is a local maximum of H. We let $t \neq e_1$ be contained in a neighborhood of e_1 to be specified later. Define $\Lambda_t := \{i : i \in [m] \setminus \{1\}, t_i > 0\}$. Then,

$$\begin{aligned} H(e_1) - H(t) &= \alpha_1 h(\beta_1^2) + \sum_{i=2}^{m} \alpha_i h(0) - \sum_{i=1}^{m} \alpha_i h(\beta_i^2 t_i) \\ &= \alpha_1 (h(\beta_1^2) - h(\beta_1^2 t_1)) - \sum_{i=2}^{m} \alpha_i (h(\beta_i^2 t_i) - h(0)) \\ &= \alpha_1 \beta_1^2 (1 - t_1) \frac{h(\beta_1^2) - h(\beta_1^2 t_1)}{\beta_1^2 (1 - t_1)} - \sum_{i \in \Lambda_t} \alpha_i \beta_i^2 t_i \frac{h(\beta_i^2 t_i) - h(0)}{\beta_i^2 t_i} \ . \end{aligned}$$

We denote the slopes given by the difference quotients in the above equation as follows:

$$m_i^{\ell}(t_i) := \frac{h(\beta_i^2 t_i) - h(0)}{\beta_i^2 t_i} \qquad\qquad m_i^{r}(t_i) := \frac{h(\beta_i^2) - h(\beta_i^2 t_i)}{\beta_i^2 (1 - t_i)} \ .$$

Thus,

$$H(e_1) - H(t) = \alpha_1 \beta_1^2 (1 - t_1) m_1^{r}(t_1) - \sum_{i \in \Lambda_t} \alpha_i \beta_i^2 t_i m_i^{\ell}(t_i) \ . \tag{7}$$

Case. $h'(0) = 0$

Let $C = m_1^r(\frac{1}{2})$. As h is strictly convex, Lemma 19 implies that $C > 0$. As $h'(0) = 0$, we have that $\lim_{x\to 0^+} \frac{h(x)-h(0)}{x} = 0$. For each $i \in [m]$ there exists $\delta_i > 0$ such that for all $x < \delta_i$, $\frac{h(x)-h(0)}{x} < C\frac{\alpha_1\beta_1^2}{\alpha_i\beta_i^2}$. We choose $t^* \in \Delta^{m-1}$ such that $t_1^* > \frac{1}{2}$ and $0 < t_i^* < \frac{\delta_i}{\beta_i^2}$ for $i \neq 1$. Fix $t \in \Delta^{m-1} \setminus \{e_1\}$ such that $|(e_1 - t)_i| < |(e_1 - t^*)_i|$ holds for each i. Then for each $i \neq 1$, $\alpha_i\beta_i^2 m_i^\ell(t_i) < \alpha_1\beta_1^2 C$ holds. Also, $m_1^r(t_1) > C$. It follows that

$$H(e_1) - H(t) > \alpha_1\beta_1^2 \left[(1 - t_1)C - \sum_{i\in\Lambda_t} t_i C\right] = 0 \ ,$$

since $\sum_{i=1}^m t_i = 1$. Thus, e_1 is a maximum of H.

Case. $h'(0) = -\infty$

Let $C = m_1^r(\frac{1}{2})$. Since $h'(0) = -\infty$, it follows that for $i \neq 1$, there exists $\delta_i > 0$ such that for any $x \in (0, \delta_i)$, $m_i^\ell(x) < \frac{\alpha_1\beta_1^2 C}{\alpha_i\beta_i^2}$ holds. Thus for any $t \in \Delta^{m-1}$ such that $t_1 \geq \frac{1}{2}$ and $t_i < \delta_i$ for each $i \neq 1$, it follows (using equation (7)) that

$$H(e_1) - H(t) > \alpha_1\beta_1^2 \left[(1 - t_1)C - \sum_{i\in\Lambda_t} t_i C\right] = 0$$

since m_1^r is a strictly increasing function and since $\sum_{i=1}^m t_i = 1$. It follows that e_1 is a maximum of H in this case as well. □

The proof of Theorem 2 now follows quite easily.

Proof of Theorem 2. Recall that we are working in the coordinate system where Z_i is given by the i^{th} canonical vector e_i. Recall also that $H : \Delta^{m-1} \to \mathbb{R}$ is defined by $H(t) := F_g \circ \psi^{-1}(t) = \frac{1}{m}\sum_{i=1}^m \alpha_i h(\beta_i^2 t_i)$ where $h(x) = g(\sqrt{x})$ is a convex function. In particular, $u \in S^{m-1}$ is a local optima of F_g if and only if $\psi(u)$ is a local optima of H.

Noting that $h_i(x) := \alpha_i h(\beta_i^2 x)$ is a strictly convex function, Lemma 7 implies that F_g has no optima outside the set $\{\pm Z_i\}_{i=1}^m$. That $\{\pm Z_i\}_{i=1}^m$ are among the optima of F_g is a consequence of Lemma 11. □

A.2 Necessity Argument for Weighted Basis Recovery

We now demonstrate the necessity conditions of Theorem 3. As this theorem involves two main parts, each piece will be considered separately.

Necessity of property P1. We will now demonstrate that the strict convexity of $t \mapsto g(\sqrt{t})$ is necessary to rule out additional local maxima for all valid choices of the weights α_i, β_i. As g is assumed to be twice differentiable, we use the differential definition of convexity.

Lemma 12. *Let $g : [0, \infty) \to \mathbb{R}$ be a continuous function such that g is twice continuously differentiable in (a, b) for some $0 < a < b$ and such that $\frac{d^2}{dt^2} g(\sqrt{t}) < 0$ for some $t \in (a, b)$. Then there exist weights $v, w > 0$ so that $F_g(x, y) = vg(|x|/\sqrt{v}) + wg(|y|/\sqrt{w})$ has a strict local maximum relative to S^1 that is not in $\{\pm e_1, \pm e_2\}$. More precisely, if $v = w = 1/(2t)$, then $(x, y) = (1/\sqrt{2}, 1/\sqrt{2})$ is a strict local maximum of F_g relative to S^1.*

Proof. We use the homeomorphism from the proof of Lemma 8. Let $h(x) = g(\sqrt{x})$. It is enough to show that $1/2$ is a strict local maximum of

$$H(x) = vh(x/v) + wh((1-x)/w)$$

in $[0,1]$. With our choice of weights v and w, we have $H'(x) = h'(x/v) - h'((1-x)/w)$ and $H'(1/2) = 0$. Similarly, $H''(x) = h''(x/v)/v + h''((1-x)/w)/w$ and $H''(1/2) = h''(t)/v + h''(t)/w < 0$. This completes the proof. □

Lemma 13. *Let $g : [0,\infty) \to \mathbb{R}$ be a continuous function such that g is twice continuously differentiable in (a,b) for some $0 < a < b$ and such that g is not strictly convex in (a,b). Then there exist weights $v, w > 0$ so that $F_g(x,y) = vg(|x|/\sqrt{v}) + wg(|y|/\sqrt{w})$ has a local maximum relative to S^1 that is not in $\{\pm e_1, \pm e_2\}$.*

Proof. If $\frac{d^2}{dt^2} g(\sqrt{t}) < 0$ for some $t \in (a,b)$, then Lemma 12 gives the desired conclusion. That is, it remains to show the conclusion under the assumption that $\frac{d^2}{dt^2} g(\sqrt{t}) \geq 0$ for $t \in (a,b)$.

We use the same notation from the proof of Lemma 12. Let $h(x) = g(\sqrt{x})$. It is enough to show that $1/2$ is a local maximum of

$$H(x) = vh(x/v) + wh((1-x)/w)$$

in $[0,1]$. We have that h is convex in (a,b) but not strictly convex. Let $y, z \in (a,b)$ and $\lambda \in (0,1)$ violate strict convexity, namely, $y < z$ and $h(t) = \lambda h(y) + (1-\lambda)h(z)$ where $t = \lambda y + (1-\lambda)z$. Convexity implies that h is an affine function in (y,z) so that h' is constant in a neighborhood of t. With the choice of weights $v = w = 1/2t$, we have $H'(x) = h'(2tx) - h'(2t(1-x))$ and therefore $H'(x) = 0$ in a neighborhood of $1/2$. This completes the proof. □

The necessity of P1 described by Theorem 3 part 1 is a direct consequence of Lemma 13.

Necessity of property P2. We continue with the definitions that $h(x) := g(\sqrt{x})$ and that $H(t) := F_g \circ \psi^{-1}(t)$ where $g : [0,\infty) \to \mathbb{R}$ is continuous and twice differentiable away from the origin as before. As previously noted, the properties P1 and P2 on g are equivalent to the properties P1* and P2* on h. As F_g takes on a strict local maximum at $Z_1 = e_1$ if and only if H takes on a strict local maximum at e_1, the necessity of P2 in Theorem 3 part 2 is an immediate consequence of the following Lemma:

Lemma 14. *Suppose that for a function $h : [0,\infty) \to \mathbb{R}$, P1* holds but P2* does not hold. Thus, the right derivative $h'(0) \notin \{0, -\infty\}$. Then there exists positive constants $\alpha_1, \ldots, \alpha_m$ and $\beta_1, \ldots, \beta_m$ such that the function $H : \Delta^{m-1} \to \mathbb{R}$ given by $H(t) = \sum_{i=1}^m \alpha_i h(\beta_i^2 t_i)$ does not take on a strict local maximum at e_1.*

Proof. From equation (7), we have that for any choice of $t \in \Delta^{m-1}$,

$$H(e_1) - H(t) = \alpha_1 \beta_1^2 (1-t_1) m_1^r(t_1) - \sum_{i \in \Lambda_t} \alpha_i \beta_i^2 t_i m_i^\ell(t_i) \ , \tag{8}$$

where the functions m_i^ℓ and m_i^r and the set Λ_t are defined in the proof of Lemma 11. We then have two cases to consider:

Case. $h'(0) > 0$

First, we note that $h'(0) < \infty$. To see this, we note that $h'(0) = \lim_{x\to 0^+} \frac{h(x)-h(0)}{x-0}$. But by Lemma 19, $x \mapsto \frac{h(x)-h(0)}{x-0}$ is an increasing function of x. Further, it is finite for any choice of $x \in (0, \infty)$.

Fix $\beta_1 = \beta_2 = \cdots = \beta_m = \beta > 0$. Let $M = D_- h(\beta^2)$. As a consequence of Lemma 19, it follows that for each $t \in \Delta^{m-1} \setminus \{e_1\}$, $m_1^r(t) < M$ holds.

Let $\alpha_1 = \frac{h'(0)}{M}$. For each $i \neq 1$, let $\alpha_i = 1$. Fix any $t \in \Delta^{m-1} \setminus \{e_1\}$. From Lemma 19, we have that $m_i^\ell(t_i) > h'(0)$ for each $i \in \Lambda_t$. It follows from equation (8) that

$$H(e_1) - H(t) < \beta^2(1-t_1)h'(0) - \sum_{i\in\Lambda_t} \beta^2 t_i h'(0) = 0 \ ,$$

making e_1 a minimum rather than a maximum of H.

Case. $h'(0) < 0$

By Lemma 21, it follows that there exists $\delta > 0$ such that for $x \in (0, \delta)$, we have that $h'(x) < \frac{1}{2}h'(0)$. Fix $\beta_1 = \cdots = \beta_m = \sqrt{\frac{1}{2}\delta}$. As a consequence of Lemma 19, it follows that for $t \in \Delta^{m-1}$, $m_1^r(t_1) < h'(\beta_i^2 t_1) < \frac{1}{2}h'(0)$. Also by Lemma 19, $m_i^\ell(t) \geq h'(0)$ for each $i \in [m]$. Fix $\alpha_1 = 2$ and $\alpha_2 = \cdots = \alpha_m = 1$. It follows from equation (8) that for any $t \in \Delta^{m-1} \setminus \{e_1\}$,

$$H(e_1) - H(t) < \frac{\delta}{2}\left[(1-t_1)h'(0) - \sum_{i\in\Lambda_t} t_i h'(0)\right] = 0$$

Thus e_1 is a minimum rather than a maximum of H. □

B Clustering via the symmetric normalized Laplacian

All arguments in the main text apply to the Graph Laplacians L and L_{rw}. However, L_{sym} is also very popular for use in spectral clustering. In this appendix, we discuss how "arbitrary" functions can be used for Spectral Clustering via L_{sym} in the setting where G consists of m connected components. The Lemmas and their proofs in this appendix are intended to highlight the differences in showing that L_{sym} is admissible for the proposed spectral algorithms in place of L or L_{rw} from the main text.

Recall that L_{sym} is defined from L via $L_{\text{sym}} = D^{-1/2} L D^{-1/2}$. Whereas taking an orthogonal basis of $\mathcal{N}(L)$ or $\mathcal{N}(L_{\text{rw}})$ produces embedded points which are orthogonal and of fixed norm within any particular class, using $\mathcal{N}(L_{\text{sym}})$ produces embedded points along perpendicular rays but with varying intra-class norms as will be seen in Lemma 15. Despite this difference, the basic arguments in the main text generalize to the L_{sym} embedding with a bit more cluttering of symbols. In particular, given a contrast function g meeting property P1 from the main text, the proposed algorithms which worked for spectral clustering using L and L_{rw} also work for spectral clustering using L_{sym}.

B.1 Structure of the null space

Let $G = (V, A)$ be a graph containing m connected components such that the i^{th} component has vertices with indices in the set $\mathcal{S}_i$. For any set $C \subset V$, we define

$$\delta(C) := \sum_{i\in C} d_{ii} \ .$$

Lemma 15. *Let G be a similarity graph consisting of m connected components for which L_{sym} is well defined. Let the vertex indices be partitioned into sets $\mathcal{S}_1, \ldots, \mathcal{S}_m$ corresponding to the m connected components. Then, $\dim(\mathcal{N}(L_{\text{sym}})) = m$. If $X = (x_{\cdot 1}, \ldots, x_{\cdot m})$ contains a scaled basis of $\mathcal{N}(L_{\text{sym}})$ in its columns such that $\|x_{\cdot i}\| = \sqrt{n}$, then there exist m mutually orthogonal unit vectors $Z_1, \ldots, Z_m$ such that whenever $i \in \mathcal{S}_j$, the row vector $x_{i\cdot} = \sqrt{n d_{ii} \delta(\mathcal{S}_j)^{-1}} Z_j^T$.*

Proof. An important property of the symmetric Laplacian (see for instance [21] Proposition 3) is that for any $u \in \mathbb{R}^n$,

$$u^T L_{\text{sym}} u = \frac{1}{2} \sum_{i,j \in V} a_{ij} \left(\frac{u_i}{d_{ii}^{1/2}} - \frac{u_j}{d_{jj}^{1/2}} \right)^2 . \tag{9}$$

L_{sym} is positive semi-definite, and u is a 0-eigenvector of L_{sym} if and only if plugging u into equation (9) yields 0. Let $v_{\mathcal{S}_j}$ be the vector such that

$$v_{\mathcal{S}_j} = \begin{cases} d_{ii}^{1/2} & \text{if } i \in \mathcal{S}_j. \\ 0 & \text{otherwise} \end{cases} .$$

Then, $B = (\delta(\mathcal{S}_1)^{-1/2} v_{\mathcal{S}_1}, \ldots, \delta(\mathcal{S}_d)^{-1/2} v_{\mathcal{S}_d})$ contains an orthonormal basis for $\mathcal{N}(L_{\text{sym}})$ in its columns.

Defining $M_{\mathcal{S}_i} = v_{\mathcal{S}_i} v_{\mathcal{S}_i}^T$, we get:

$$P_{\mathcal{N}(L)} = BB^T = \sum_{i=1}^{m} \delta(\mathcal{S}_i)^{-1} M_{\mathcal{S}_i} . \tag{10}$$

But the projection matrix can be constructed from any orthonormal basis of $\mathcal{N}(L)$. In particular, $P_{\mathcal{N}(X)} = \frac{1}{n} X X^T$ as well. Hence, $\frac{1}{n} \langle x_{i\cdot}, x_{j\cdot} \rangle = (P_{\mathcal{N}(L)})_{ij} = \delta(\mathcal{S}_\ell)^{-1} d_{ii}^{1/2} d_{jj}^{1/2}$ precisely when there exists $\ell \in [m]$ such that $i, j \in \mathcal{S}_\ell$. Otherwise, $x_{i\cdot} \perp x_{j\cdot}$.

Note that for $i, j \in \mathcal{S}_\ell$,

$$\cos(\angle(x_{i\cdot}, x_{j,\cdot})) = \frac{\langle x_{i\cdot}, x_{j\cdot} \rangle}{\langle x_{i\cdot}, x_{i\cdot} \rangle^{1/2} \langle x_{j\cdot}, x_{j\cdot} \rangle^{1/2}} = \frac{n \delta(\mathcal{S}_\ell)^{-1} d_{ii}^{1/2} d_{jj}^{1/2}}{n^{1/2} \delta(\mathcal{S}_\ell)^{-1/2} d_{ii}^{1/2} n^{1/2} \delta(\mathcal{S}_\ell)^{-1/2} d_{jj}^{1/2}} = 1$$

giving that points from the same cluster lie on the same ray from the origin. It follows that there are m mutually orthogonal unit vectors, $Z_1, \ldots, Z_m$ such that $x_{i\cdot} = \sqrt{n d_{ii} \delta(\mathcal{S}_\ell)^{-1}} Z_\ell^T$ for each $i \in \mathcal{S}_\ell$. □

B.2 Contrast admissibility under the symmetric Laplacian

Let G be a similarity graph containing m connected components, and let L_{sym} be constructed from G. Then, let X contain the scaled eigenvecotors of L_{sym} constructed in accordance with Lemma 15. We let $Z_1, \ldots, Z_m$ denote the directional vectors of the same name from Lemma 15. Parallel to the main text, we define functions $F_g : S^{m-1} \to \mathbb{R}$ using an "arbitrary" continuous function $g : [0, \infty) \to \mathbb{R}$ such that

$$F_g(u) := \frac{1}{n} \sum_{j=1}^{n} g(|\langle u, x_{j\cdot} \rangle|) = \frac{1}{n} \sum_{i=1}^{m} \sum_{j \in \mathcal{S}_i} g\left(|\langle u, Z_i \rangle| \|x_{j\cdot}\| \right) . \tag{11}$$

We assume only that property P1 from the main text holds for g. Then, demonstrating that F_g has no extraneous local maxima is a relatively simple generalization of the results from the main text.

Lemma 16. *Suppose that $g : [0,\infty) \to \mathbb{R}$ is continuous and satisfies P1. Let F_g be constructed from g according to equation (11). Then, the set of local maxima of F_g is contained in $\{\pm Z_i : i \in [m]\}$.*

Proof. Define $g_i : [0,\infty) \to \mathbb{R}$ by $g_i(t) = \frac{1}{n}\sum_{j\in\mathcal{S}_i} g(t\|x_{i\cdot}\|)$. Then, it follows from equation (11) that $F_g(u) = \sum_{i=1}^m g_i(|\langle u, Z_i\rangle|)$. Since $t \mapsto g(\sqrt{t})$ is strictly convex, $t \mapsto g_i(\sqrt{t})$ is strictly convex for each g_i. Lemma 8 implies (letting each Z_i take on the role of e_i) that F_g has no local maxima outside the set $\{\pm Z_i : i \in [m]\}$. □

What remains to be seen is that the directions $\{\pm Z_i\}_{i=1}^m$ are local maxima of F_g. As before, we identify $Z_1, \ldots, Z_m$ with the canonical directions $e_1, \ldots, e_m$ in an unknown coordinate system, and we use that the simplex $\Delta^{m-1} := \mathrm{conv}(e_1, \ldots, e_m)$ is homeomorphic to $Q_1 := S^{m-1} \cap [0,\infty)^m$ under the map $\psi : Q_1 \to \Delta^{m-1}$ defined by $\psi_i(u) = u_i^2$.

Lemma 17. *Let $h : [0,\infty) \to \mathbb{R}$ be a strictly convex function. Let $H : \Delta^{m-1} \to \mathbb{R}$ be given by $H(u) = \frac{1}{n}\sum_{i=1}^m \sum_{j\in\mathcal{S}_i} h(u_i\|x_{j\cdot}\|^2)$. Then the set $\{e_i\}_{i=1}^m$ is contained in the set of strict local maxima of H.*

Proof. By the symmetries of H, it suffices to show that e_1 is a strict local maximum of H. To see this, choose $u \neq e_1$ from a neighborhood of e_1 relative to Δ^{m-1} to be specified later. Let $\Lambda_u = \{i : i \in [m] \setminus \{1\}, u_i \neq 0\}$. Then,

$$\begin{aligned}
&H(e_1) - H(u) \\
&= \frac{1}{n}\left[\sum_{j\in\mathcal{S}_1} h(\|x_{j\cdot}\|^2) + \sum_{i=2}^m \sum_{j\in\mathcal{S}_i} h(0) - \sum_{i=1}^m \sum_{j\in\mathcal{S}_i} h(u_i\|x_{j\cdot}\|^2)\right] \\
&= \frac{1}{n}\left[\sum_{j\in\mathcal{S}_1} \big(h(\|x_{j\cdot}\|^2) - h(u_1\|x_{j\cdot}^2\|)\big) - \sum_{i=2}^m \sum_{j\in\mathcal{S}_i} \big(h(u_i\|x_{j\cdot}\|^2) - h(0)\big)\right] \\
&= \frac{1}{n}\left[\sum_{j\in\mathcal{S}_1} \|x_{j\cdot}\|^2(1-u_1)\frac{h(\|x_{j\cdot}\|^2) - h(u_1\|x_{j\cdot}^2\|)}{\|x_{j\cdot}\|^2(1-u_1)} - \sum_{i\in\Lambda_u}\sum_{j\in\mathcal{S}_i} u_i\|x_{j\cdot}\|^2 \frac{h(u_i\|x_{j\cdot}\|^2) - h(0)}{u_i\|x_{j\cdot}\|^2}\right] .
\end{aligned}$$

We have written $H(e_1) - H(u)$ as a weighted sum of difference quotients (slopes). We would like to apply Lemma 20 in order to demonstrate that there is a neighborhood B of e_1 relative to Δ^{m-1} such that $u \in B \setminus \{e_1\}$ implies $H(e_1) - H(u) < 0$. First, we notice that for each $x_{j\cdot}$, u breaks the interval left and right pieces, yielding two slopes of interest:

$$m_{ij}^\ell = \frac{h(u_i\|x_{j\cdot}\|^2) - h(0)}{u_i\|x_{j\cdot}\|^2} \qquad \text{and} \qquad m_{ij}^r = \frac{h(\|x_{j\cdot}\|^2) - h(u_i\|x_{j\cdot}\|^2)}{\|x_{j\cdot}\|^2(1-u_i)} .$$

Thus,

$$H(e_1) - H(u) = \frac{1}{n}\left[\sum_{j\in S_1} \|x_{j\cdot}\|^2(1-u_1)m_{1j}^r - \sum_{i\in\Lambda_u}\sum_{j\in S_i} u_i\|x_{j\cdot}\|^2 m_{ij}^\ell\right] .$$

Let $B = \{u : u_i < \frac{\min_j\|x_{j\cdot}\|^2}{\max_j\|x_{j\cdot}\|^2} \text{ for all } i \neq 1\}$. Then, fixing $u \in B$ and $i \neq 1$, we have that $u_i\|x_{j_1\cdot}\|^2 < \|x_{j_2\cdot}\|^2$ for any $j_1 \in \mathcal{S}_i$ and $j_2 \in \mathcal{S}_1$. Let $m_{\max}^\ell := \max\{m_{ij}^\ell : i \in \Lambda_u, j \in \mathcal{S}_i\}$ and

$m^r_{\min} := \min\{m^r_{1j} : j \in \mathcal{S}_1\}$. From Lemma 20, it follows that $m^\ell_{\max} < m^r_{\min}$. Thus,

$$H(e_1) - H(u) \geq \frac{1}{n}\left[\sum_{j\in S_1} \|x_{j\cdot}\|^2(1-u_1)m^r_{\min} - \sum_{i\in\Lambda_u}\sum_{j\in S_i} u_i\|x_{j\cdot}\|^2 m^\ell_{\max}\right]$$
$$= (1-u_1)m^r_{\min} - \sum_{i=2}^{m} u_i m^\ell_{\max} = (1-u_1)[m^r_{\min} - m^\ell_{\min}] > 0$$

Thus, e_1 is a local maximum of H. □

This brings us to our main theorem, restated with F_g constructed from the embedded points of L_{sym}.

Theorem 18. *Let $g : [0,\infty) \to \mathbb{R}$ be a continuous function satisfying P1. If F_g is defined from g according to equation (11), then $\{\pm Z_i : i \in [m]\}$ is a complete enumeration of the local maxima of F_g.*

Proof. Let Λ denote the set of local maxima of F_g. That $\Lambda \subset \{\pm Z_i : i \in [m]\}$ is immediate from Lemma 16. To see that $\Lambda \supset \{\pm Z_i : i \in [m]\}$, we note that there is a natural mapping between Δ^{m-1} and a quadrant of S^{m-1}.

The set $\{\pm Z_i : i \in [m]\}$ gives an unknown, orthonormal basis of our space. We may without loss of generality work in the coordinate system where $e_1, \ldots, e_m$ coincide with $Z_1, \ldots, Z_m$. Let $Q_1 = S^{m-1} \cap [0,\infty)^{m-1}$ give the first quadrant of the unit sphere. By the symmetries of the problem, it suffices to show that $\{e_1, \ldots, e_m\}$ are maxima of F_g. However, the map $\psi : Q_1 \to \Delta^{m-1}$ defined by $(\psi(u))_i = u_i^2$ is a homeomorphism. Defining $H : \Delta^{m-1} \to \mathbb{R}$ by $H(t) = F_g(\psi^{-1}(t))$, then $t \in \Delta^{m-1}$ is a local maximum of H if and only if $\psi^{-1}(t)$ is a local maximum of F_g relative to Q_1.

Note that $H(t) = \frac{1}{n}\sum_{i=1}^m \sum_{j\in\mathcal{S}_i} g(\sqrt{t_i}\|x_{j\cdot}\|^2)$. As $y \mapsto g(\sqrt{y})$ is convex, it follows by Lemma 17 that $\{e_i\}_{i=1}^m$ are local maxima of H. Hence, using the symmetries of F_g, $\{\pm Z_i : i \in [m]\} \supset \Lambda$. □

C Facts about convex functions

In this section, intervals can be open, half open, or closed.

There is a large literature studying the properties of convex functions. As strict convexity is considered more special than convexity, results are typically stated in terms of convex functions. The following characterization of strict convexity is a version of Proposition 1.1.4 of [8] for strictly convex functions, and can be proven in a similar fashion.

Lemma 19. *For an interval I, let $f : I \to \mathbb{R}$ be a strictly convex function. Then, fixing any $x_0 \in I$, the slope function defined by $m(x) := \frac{f(x)-f(x_0)}{x-x_0}$ is strictly increasing on $I \setminus \{x_0\}$.*

The following result is largely a consequence of Lemma 19.

Lemma 20. *Let I be an interval and let $f : I \to \mathbb{R}$ be a convex function. Suppose that $(a,b) \subset I$ and $(c,d) \subset I$ are such that $a \leq c$ and $b \leq d$ with at least one of the inequalities being strict. Then,*

$$\frac{f(b)-f(a)}{b-a} < \frac{f(d)-f(c)}{d-c}$$

Proof. If $c = a$, then $\frac{f(d)-f(a)}{d-a} = \frac{f(d)-f(c)}{d-c}$ trivially. Otherwise, $a < c$, and by Lemma 19, we have that $\frac{f(d)-f(a)}{d-a} < \frac{f(d)-f(c)}{d-c}$ By similar reasoning, $\frac{f(b)-f(a)}{b-a} \leq \frac{f(d)-f(a)}{d-a}$ (with equality if and only if $d = b$). As by assumption, $a = b$ and $c = d$ cannot both hold, it follows that $\frac{f(b)-f(a)}{b-a} \leq \frac{f(d)-f(a)}{d-a} \leq \frac{f(d)-f(c)}{d-c}$ with at least one of the inequalities being strict. □

The following result can be found for instance in Remark 4.2.2 of [8]

Lemma 21. *Given an interval I and a function $f : I \to \mathbb{R}$, then the left derivative D_-f is left-continuous and the right derivative D_+f is right-continuous respectively whenever they are defined (that is, finite).*

D Perturbation of the Spectral Embedding

In this appendix and in the subsequent appendix E, we demonstrate that even under a perturbation from the setting in which our similarity graph G contains m-connected components, the maxima structure of the F_g used in our algorithms HBROPT and HBRENUM is approximately preserved. In this appendix, we demonstrate that under a sufficiently small perturbation, the configuration of points resulting from the Laplacian embedding in Proposition 4 and Lemma 15 are approximately preserved. In appendix E, we will further demonstrate that for certain choices of contrast functions that the local maxima structure of F_g constructed from the embedded data are also approximately preserved.

Throughout this section, we assume that $\mathcal{L}$ is a graph Laplacian (the reader's choice among L or L_{sym}) for a graph G consisting of m connected components. We exclude the case of L_{rw} since it is not symmetric and would make the perturbation argument slightly more complicated, though similar results can be obtained with L_{rw}. Letting $\lambda_1(\mathcal{L}) \leq \lambda_2(\mathcal{L}) \leq \cdots \leq \lambda_n(\mathcal{L})$ denote the eigenvalues of $\mathcal{L}$, then by our assumptions it follows that $0 = \lambda_1(\mathcal{L}) = \lambda_2(\mathcal{L}) = \cdots = \lambda_m(\mathcal{L}) < \lambda_{m+1}(\mathcal{L})$. We will denote by $\delta(\mathcal{L}) = \lambda_{m+1}(\mathcal{L}) - \lambda_m(\mathcal{L})$ the eigengap between $\mathcal{N}(\mathcal{L})$ and the range of $\mathcal{L}$ (denoted as $\mathcal{R}(\mathcal{L})$). We further assume that $\tilde{\mathcal{L}} = \mathcal{L} + H$ is a perturbation of $\mathcal{L}$ constructed from a graph whose similarity matrix is a perturbation of a graph G consisting of m connected components. In particular H is a symmetric matrix, making $\tilde{\mathcal{L}}$ symmetric, and $\|H\|$ is viewed as being small. Our main result in this section is the following.

Theorem 22. *Suppose that $\tilde{X} = (\tilde{x}_{\cdot 1}, \ldots, \tilde{x}_{\cdot m})$ contains the scaled lowest m eigenvectors of $\tilde{\mathcal{L}}$ such that each $\|\tilde{x}_{i\cdot}\| = \frac{1}{\sqrt{n}}$. Then there exists $X = (x_{\cdot 1}, \ldots, x_{\cdot m})$ with columns forming a scaled orthogonal basis of $\mathcal{N}(\mathcal{L})$ such that each $\|x_{\cdot i}\| = \frac{1}{\sqrt{n}}$ and $\frac{1}{\sqrt{n}}\|X - \tilde{X}\| \leq \frac{2\|H\|}{\delta(\mathcal{L}) - \|H\|}$. In particular, we have the following perturbation bounds on the embedding:*

1. *If $v \in \mathbb{R}^m$ is a unit vector, then $\frac{1}{n}\sum_{i=1}^n (\langle \tilde{x}_{i\cdot}, v\rangle - \langle x_{i\cdot}, v\rangle)^2 \leq \frac{4\|H\|^2}{(\delta(\mathcal{L}) - \|H\|)^2}$.*
2. *$\frac{1}{n}\sum_{i=1}^n \|\tilde{x}_{i\cdot} - x_{i\cdot}\|^2 \leq \frac{4\|H\|^2 m}{(\delta(\mathcal{L}) - \|H\|)^2}$.*

In particular, the orthogonal basis structure from Proposition 4 which arises when $\mathcal{L} = L$ (or the orthogonal ray structure for $\mathcal{L} = L_{\text{sym}}$ from Lemma 15) is approximately preserved.

We note that a related result for k-means was provided by [16] when embedding with L_{sym} and unit normalizing the embedded data.

We now proceed with the proof of Theorem 22. The following lemma is a direct implication of the Davis-Kahan $\sin\Theta$ theorem [5].

Lemma 23. *Suppose that $X = (x_{\cdot 1}, \ldots, x_{\cdot m})$ is an orthogonal basis of $\mathcal{N}(\mathcal{L})$ and that $\tilde{X} = (x_{\cdot 1}, \ldots, x_{\cdot m})$ are the bottom m eigenvectors of $\tilde{\mathcal{L}}$. Then, $\|P_{\mathcal{N}(X)}P_{\mathcal{R}(\tilde{X})}\| \leq \frac{\|H\|}{\delta(\mathcal{L}) - \|H\|}$.*

We now decompose the bound from Lemma 23 in order to write a bound in terms of the actual entries of X and $\tilde{X}$. This is complicated by the fact that X is only meaningfully defined up to a rotation of its columns within the subspace $\mathcal{N}(\mathcal{L})$. In Lemma 23, this is reflected by the fact that we are bounding the change in the projection operators onto the subspaces spanned by X

and $\tilde{X}$ rather than directly bounding the change in the eigenvectors themselves. Indeed, if X and Y contained two different (scaled) bases of $\mathcal{N}(\mathcal{L})$ in their columns, then $R = \frac{1}{n}X^TY$ would be a rotation matrix providing a transition between these basis systems, making $Y = XR$. We may then think of $\frac{1}{n}X^T\tilde{X}$ as being an approximate rotation matrix. The following Lemmas provide the key steps in creating a bound which takes into account the missing rotation explicitly.

Lemma 24. *Suppose that $X = (x_{\cdot 1}, \ldots, x_{\cdot m})$ contains a scaled orthogonal basis of $\tilde{\mathcal{L}}$ such that each $\|x_{\cdot i}\| = \frac{1}{\sqrt{n}}$ Suppose that $\tilde{X} = (\tilde{x}_{\cdot 1}, \ldots, \tilde{x}_{\cdot m})$ contains the scaled lowest m eigenvectors of $\tilde{\mathcal{L}}$ such that each $\|\tilde{x}_{\cdot i}\| = \frac{1}{\sqrt{n}}$. Then, for any rotation matrix $R \in \mathbb{R}^{m\times m}$, $\frac{1}{\sqrt{n}}\|XR - \tilde{X}\| \leq \frac{\|H\|}{\delta(\mathcal{L})-\|H\|} + \|\frac{1}{n}\tilde{X}^TX - R\|$.*

Proof. Applying Lemma 23, we obtain have that $\|P_{\mathcal{N}(X)}P_{\mathcal{R}(\tilde{X})}\| \leq \frac{\|H\|}{\delta(\mathcal{L})-\|H\|}$. Since $\frac{1}{\sqrt{n}}\tilde{X}^T$ treated as a map from $\mathcal{R}(\tilde{X})$ to $\mathbb{R}^m$ is an isometry, we obtain that

$$\|P_{\mathcal{N}(X)}P_{\mathcal{R}(\tilde{X})}\| = \|(I - \frac{1}{n}XX^T)\frac{1}{n}\tilde{X}\tilde{X}^T\| = \|(I - \frac{1}{n}XX^T)\frac{1}{\sqrt{n}}\tilde{X}\| = \frac{1}{\sqrt{n}}\|\tilde{X} - \frac{1}{n}XX^T\tilde{X}\| \ .$$

We fix a rotation matrix R and expand $\frac{1}{n}X^T\tilde{X} = \frac{1}{n}X^T\tilde{X} - R + R$ to obtain:

$$\begin{aligned}\|P_{\mathcal{N}(X)}P_{\mathcal{R}(\tilde{X})}\| &= \frac{1}{\sqrt{n}}\|\tilde{X} - XR + X(R - \frac{1}{n}X^T\tilde{X})\| \\ &\geq \frac{1}{\sqrt{n}}(\|\tilde{X} - XR\| - \|X(R - \frac{1}{n}X^T\tilde{X}\|) \ . \end{aligned} \tag{12}$$

Treating $\frac{1}{\sqrt{n}}X$ as a linear map from $\mathcal{R}(X)$ to $\mathbb{R}^m$ defined by $v \mapsto \frac{1}{\sqrt{n}}v^TX$, then $\frac{1}{\sqrt{n}}X$ is an isometry. As such, $\|\frac{1}{\sqrt{n}}X(R - \frac{1}{n}X^T\tilde{X}\| = \|R - \frac{1}{n}X^T\tilde{X}\|$. Rearranging terms in equation (12) yields:

$$\frac{1}{\sqrt{n}}\|\tilde{X} - XR\| \leq \|P_{\mathcal{N}(X)}P_{\mathcal{R}(\tilde{X})}\| + \|R - \frac{1}{n}X^T\tilde{X}\| \leq \frac{\|H\|}{\delta(\mathcal{L}) - \|H\|} + \|R - \frac{1}{n}X^T\tilde{X}\| \ . \qquad \square$$

Lemma 25. *Suppose that X and $\tilde{X}$ are as in Lemma 24. Then, there exists a rotation matrix $R \in \mathbb{R}^{m\times m}$ such that $\|R - \frac{1}{n}X^T\tilde{X}\| \leq \|P_{\mathcal{N}(X)}P_{\mathcal{R}(\tilde{X})}\| \leq \frac{\|H\|}{\delta(\mathcal{L})-\|H\|}$.*

Proof. Letting $U\Sigma V^T$ be the singular value decomposition of $\frac{1}{n}X^T\tilde{X}$, we consider the rotation $R = UV^T$. For this choice of R, $\|R - \frac{1}{n}X^T\tilde{X}\| = \|U(I - \Sigma)V^T\| = \|I - \Sigma\|$. It suffices to show that all of the singular values of $\frac{1}{n}X^T\tilde{X}$ are contained in the interval $[1 - \|P_{\mathcal{N}(X)}P_{\mathcal{R}(\tilde{X})}\|, 1]$.

We denote the singular values of $\frac{1}{n}X^T\tilde{X}$ by the decreasing sequence $\sigma_1(X^T\tilde{X}) \geq \cdots \geq \sigma_m(X^T\tilde{X})$. An upper bound is given by $\sigma_1(X^T\tilde{X}) \leq \|\frac{1}{n}X^T\tilde{X}\| \leq \frac{1}{n}\|X^T\|\|\tilde{X}\| \leq 1$.

We now find a lower bound on the singular values. Construct the matrix $Y \in \mathbb{R}^{n\times(n-m)}$ such that the columns of $[X\ Y]$ forms a scaled orthogonal basis of $\mathbb{R}^n$ with each column having $\frac{1}{\sqrt{n}}$ norm. By construction, $\sigma_1(\frac{1}{n}[X\ Y]^T\tilde{X}) = \cdots = \sigma_m(\frac{1}{n}[X\ Y]^T\tilde{X}) = 1$ and $\sigma_{m+1}(\frac{1}{n}[X\ Y]^T\tilde{X}) = \cdots = \sigma_n(\frac{1}{n}[X\ Y]^T\tilde{X}) = 0$. Expanding $\frac{1}{n}[X\ Y]^T\tilde{X} = \frac{1}{n}X^T\tilde{X} + \frac{1}{n}Y^T\tilde{Y}$, we obtain from a Weyl's inequality like bound for singular values [9, Theorem 3.3.16] that $\sigma_m(\frac{1}{n}[X\ Y]^T\tilde{X}) \leq \sigma_m(\frac{1}{n}X^T\tilde{X}) + \sigma_1(\frac{1}{n}Y^T\tilde{X})$. Hence, $\sigma_m(\frac{1}{n}X^T\tilde{X}) \geq \sigma_m(\frac{1}{n}[X\ Y]^T\tilde{X}) - \sigma_1(\frac{1}{n}Y^T\tilde{X}) = 1 - \|\frac{1}{n}Y^T\tilde{X}\|$. Noting that the maps $\frac{1}{\sqrt{n}}Y : \mathcal{R}(Y) \to \mathbb{R}^m$ defined by $v \mapsto 1\sqrt{n}v^TY$ and $\frac{1}{\sqrt{n}}\tilde{X}^T : \mathcal{R}(\tilde{X}) \to \mathbb{R}^m$ defined by $v \mapsto \frac{1}{\sqrt{n}}\tilde{X}^Tv$ are isometries, it follows that $\|\frac{1}{n}Y^TX\| = \|\frac{1}{n^2}YY^T\tilde{X}\tilde{X}^T\| = \|P_{\mathcal{N}(X)}P_{\mathcal{R}(\tilde{X})}\|$. In particular, $\sigma_m(\frac{1}{n}X^T\tilde{X}) \geq 1 - \|P_{\mathcal{N}(X)}P_{\mathcal{R}(\tilde{X})}\|$. $\square$

Lemmas 24 and 25 combine to provide the following bound on the embedding error.

Proposition 26. *Let $X = (x_{\cdot 1}, \ldots, x_{\cdot m})$ contain a scaled orthogonal basis of $\mathcal{N}(\mathcal{L})$ such that each $\|x_{\cdot i}\| = \frac{1}{\sqrt{n}}$, and let $\tilde{X} = (\tilde{x}_{\cdot 1}, \ldots, \tilde{x}_{\cdot m})$ be the scaled lowest m eigenvectors of $\tilde{\mathcal{L}}$ such that each $\|\tilde{x}_{\cdot i}\| = \frac{1}{\sqrt{n}}$. There exists an orthogonal matrix $R \in \mathbb{R}^{m \times m}$ such that $\frac{1}{\sqrt{n}}\|XR - \tilde{X}\| \leq \frac{2\|H\|}{\delta(\mathcal{L}) - \|H\|}$.*

Theorem 22 follows by setting X in Theorem 22 to be XR from Proposition 26.

E Robust maxima structure of the basis encoding

We have seen in theorems 2 and 3 that the maxima of F_g constructed from a contrast g satisfying assumptions P1 and P2 are precisely the encoded basis directions $\{\pm Z_i : i \in [m]\}$. However, in practice we only expect to have access to estimates of F_g and it derivatives. In this section, we demonstrate that F_g has a robust maxima structure under some natural robustness assumptions. In particular, given strictly positive constants $\alpha_i, \beta_i, c_{\min}, c_{\max}$, and D, we say that the contrast g (or alternatively F_g) is $(c_{\min}, c_{\max}, D)$-robust if $F_g(u) := \sum_{i=1}^m \alpha_i g(\beta_i |\langle u, Z_i \rangle|)$ satisfies P2 and the following robust version of P1 for each $i \in [m]$:

R1. For every $x \in \mathbb{R}$, $|\frac{d^2}{dt^2} g(\sqrt{t})|_{t=x}| \in [c_{\min}, c_{\max}]$.

R2. The function $g(\sqrt{|t|}$ is three times differentiable, and $\left|\frac{d^3}{dt^3} g(t)|_{t=x}\right| < D$ for every $x \in \mathbb{R}$.

A consequence of R1 and P2 is that $\frac{d}{dx}(g(\sqrt{x}))|_{x=0^+} = 0$. At times we will only need assumption R1 but not assumption R2. In these cases, we will refer to the function F_g constructed from such a strongly convex g as being $c_{\min}$-$c_{\max}$ robust. For such a choice of contrast function, we will demonstrate that the maxima structure of F_g is robust to a perturbation. We will first see this for general basis encodings, and we will later see this for those functions F_g which arise from the Laplacian embedding under a perturbation from the setting in which the similarity graph G consists of m connected components.

E.1 Tangent space of the sphere

We will analyze the first and second derivative conditions for critical points on the sphere S^{m-1} while treating S^{m-1} as a differential manifold using the projective coordinates. Given a vector $v \in S^{m-1}$, we define a coordinate chart on the hemisphere $S_v^+ := \{u \in S^{m-1} : \langle u, v \rangle > 0\}$ as follows: We fix $p_1, \ldots, p_{m-1}, p_m = v$ an orthonormal basis of $\mathbb{R}^m$, and we define $\pi_v : S_v^+ \to \mathbb{R}^{m-1}$ by $(\pi_v)_i(u) = \frac{\langle u, p_i \rangle}{\langle u, v \rangle}$. This map is a bijection with inverse given by $\pi_v^{-1}(x) = \frac{1}{\|(1,\ x^T)\|} v + \sum_{i=1}^{m-1} \frac{x_i}{\|(1,\ x^T)\|} p_i$.

We note that the map π_v has a nice geometric interpretation. The range $\mathcal{R}(\pi_v) = \mathbb{R}^{m-1}$ can be viewed as a local coordinate system on plane $v + v^\perp$ tangent to S^{m-1} at the point v with v treated as the origin. In particular, we say that $x \in \mathcal{R}(\pi_v)$ is represented in $v + v^\perp$ by $v + \sum_{i=1}^{m-1} x_i p_i$, and we write this as either $x \simeq v + \sum_{i=1}^{m-1} x_i p_i$ or $\text{rep}(x) = v + \sum_{i=1}^{m-1} x_i p_i$. We note that the map $\text{rep} : \mathcal{R}(\pi_v) \to (v + v^\perp)$ is a bijection. Under this interpretation, π_v and its inverse take on very natural forms of radial projections onto the tangent plane and back onto the sphere respectively: $\pi_v(u) \simeq \frac{u}{\langle u, v \rangle}$ and $\pi_v^{-1}(x) = \frac{\text{rep}(x)}{\|\text{rep}(x)\|}$. We will sometimes refer to $T_v S^{n-1}$ the tangent space a v. $T_v S^{n-1}$ is the tangent plane $v + v^\perp$ with v being treated as the origin of the space, i.e., $T_v = v^\perp$.

As π_v is a diffeomorphism, it is clear that the optima structure of $F_g|_{S^{m-1}}$ at a point $v \in S^{m-1}$ is precisely the same as the optima structure of $F_g \circ \pi_v^{-1}$ at $\pi_v(v) = 0$. As such, it will suffice to analyze the structure of the functions $F_g \circ \pi_v^{-1}$ at various points on the sphere. We now provide formulas for the first two derivatives of $F_g \circ \pi_v$ evaluated at the point v which may be used in any second derivative test of extrema are provided in the following Proposition. For completeness, the derivation of these derivative formulas is provided in Appendix F.

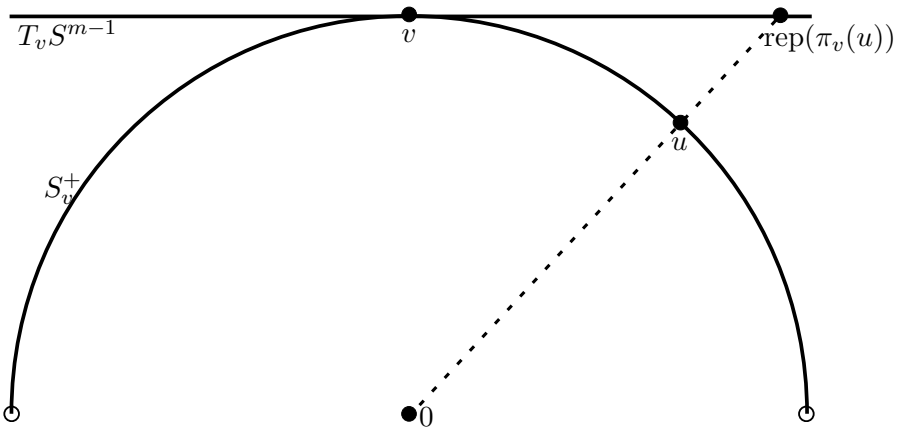

Figure 3: Mapping onto the projective space T_vS^{m-1} is a form of radial projection.

Fact 27. *Let $f : \mathbb{R}^m \to \mathbb{R}$ be twice continuously differentiable. For any $x, y \in \mathcal{R}(\pi_v)$, the first two derivatives of $f \circ \pi_v^{-1}$ evaluated at v can be interpretted as operators on T_vS^{m-1}. Letting $\xi = \mathrm{rep}(x) - v$ and $\zeta = \mathrm{rep}(y) - v$, then*

$$\langle \nabla(f \circ \pi_v^{-1})(\pi_v(v)), x\rangle = \langle \nabla f(v), \xi\rangle$$
$$x^T[\mathcal{H}(f \circ \pi_v^{-1})(\pi_v(v))]y = \xi^T[\mathcal{H}f(v) - \langle \nabla f(v), v\rangle I]\zeta$$

E.2 Robust derivative structure

In the remainder of this section, we will represent F_g in terms of the hidden basis $Z_1, \ldots, Z_m$. In particular, we define functions $h_i : \mathbb{R} \to \mathbb{R}$ by $h_i(t) := \alpha_i g(\beta_i\sqrt{|t|})$, making $F_g(u) = \sum_{i=1}^m h_i(\langle u, Z_i\rangle^2)$. We will further assume without loss of generality that $Z_1, \ldots, Z_m$ is the canonical basis, further simplifying our notation and making $F_g(u) = \sum_{i=1}^m h_i(u_i^2)$.

In this subsection, we bound where the local maxima of a perturbation of F_g may arise. More precisely, we demonstrate that for a small perturbation of $\hat{F}_g$, the maxima of $\hat{F}_g$ over the sphere include locations near the basis directions $\pm Z_1, \ldots, \pm Z_m$, and that no new spurious maxima arise. We first demonstrate the approximate stability of maxima in the following Lemma.

We denote by $\alpha_{\max} := \max_{i\in[m]} \alpha_i$, $\alpha_{\min} := \min_{i\in[m]} \alpha_i$, $\beta_{\max} := \max_{i\in[m]} \beta_i$, and $\beta_{\min} := \min_{i\in[m]} \beta_i$ in the following.

Lemma 28. *Suppose that F_g is $c_{\min}$-$c_{\max}$ robust, and that $\tilde{F}_g$ is an approximation of F_g such that for a fixed $\epsilon > 0$, $|\tilde{F}_g(u) - F_g(u)| < \epsilon$ for every $u \in S^{n-1}$. If $\epsilon < \frac{\alpha_{\min}^2\beta_{\min}^8 c_{\min}^2}{64\alpha_{\max}\beta_{\max}^4 c_{\max}}$, then for each $j \in [m]$ and each $s \in \{\pm 1\}$, there exists a local maximum v of $\tilde{F}_g$ with respect to S^{m-1} such that $\|sZ_j - v\| \le \sqrt{\frac{16\epsilon}{\alpha_{\min}\beta_{\min}^4 c_{\min}}}$.*

In order to prove Lemma 28, we make use of the following technical result:

Lemma 29. *Let $a, b \in [0, \infty)$. Then, $h_i(b) - h_i(a) \in \frac{1}{2}\alpha_i\beta_i^2(b^2 - a^2)[c_{\min},\ c_{\max}]$.*

Proof. We first expand $h_i''(t) = \frac{\mathrm{d}^2}{\mathrm{d}t^2}[\alpha_i g(\sqrt{\beta_i^2 t})]$. We define $h(t) := g(\sqrt{t})$. Then,

$$h_i''(t) = \frac{\mathrm{d}^2}{\mathrm{d}t^2}[\alpha_i h(\beta_i^2 t)] = \alpha_i\beta_i^4 h_i''(\beta_i^2 t) = \alpha_i\beta_i^4 \frac{\mathrm{d}^2}{\mathrm{d}x^2} g(\sqrt{x})|_{x=\beta_i^2 t}\ . \tag{13}$$

Using this, we obtain:

$$h_i(b) - h_i(a) = \int_a^b h_i'(x)\,\mathrm{d}x = \int_a^b \int_0^x h_i''(y)\,\mathrm{d}y\mathrm{d}x = \int_a^b \int_0^x \alpha_i\beta_i^4 \frac{\mathrm{d}^2}{\mathrm{d}t^2} g(\sqrt{t})|_{t=\beta_i^2 y}\,\mathrm{d}y\mathrm{d}x$$
$$\in \alpha_i\beta_i^4 \left[\int_a^b \int_0^x c_{\min}\,\mathrm{d}y\mathrm{d}x, \int_a^b \int_0^x c_{\max}\,\mathrm{d}y\mathrm{d}x\right] = \frac{1}{2}\alpha_i\beta_i^4(b^2-a^2)[c_{\min}, c_{\max}]\ . \qquad \square$$

Proof of Lemma 28. By the symmetries of the problem, it is sufficient to prove the case $j = 1$ and $s = +1$.

We let $\eta \in (0, \frac{1}{2}]$ be arbitrary (to be chosen later), and we consider a vector $u \in S^{m-1}$ such that $\|Z_1 - u\| = \eta$. We have:

$$\tilde{F}_g(Z_1) - \tilde{F}_g(u) = \tilde{F}_g(Z_1) - F_g(Z_1) + h_1(1) + \sum_{i=2}^m h_i(0) - \left[\tilde{F}_g(u) - F_g(u) + \sum_{i=1}^m h_i(u_i^2)\right]$$
$$\geq (h_1(1) - h_1(u_1^2)) - \sum_{i=2}^m [h_i(u_i^2) - h_i(0)] - 2\epsilon$$
$$\geq \frac{1}{2}\alpha_{\min}\beta_{\min}^4(1-u_1^4)c_{\min} - \frac{1}{2}\alpha_{\max}\beta_{\max}^4 \sum_{i=2}^m u_i^4 c_{\max} - 2\epsilon\ . \tag{14}$$

We note that since $\|Z_1 - u\|^2 = \eta^2$, we obtain:

$$\eta^2 = (1-u_1)^2 + \sum_{i=2}^m u_i^2 = (1 - 2u_1 + u_1^2) + (1 - u_1^2) = 2(1-u_1)$$

In particular, $u_1 = 1 - \frac{1}{2}\eta^2$ and $u_1^2 = 1 - \frac{1}{2}\eta^2 + \frac{1}{4}\eta^4$. By expanding $[1+u_1^2][1-u_1^2]$ in terms of η and using that $\eta \leq \frac{1}{2}$, we see that

$$1 - u_1^4 = (1+u_1^2)(1-u_1^2) = [2 - \frac{1}{2}\eta^2 + \frac{1}{4}\eta^4][\frac{1}{2}\eta^2 - \frac{1}{4}\eta^4] = \eta^2 - \frac{3}{4}\eta^4 + \frac{1}{4}\eta^8 - \frac{1}{16}\eta^{16} > \frac{1}{2}\eta^2\ .$$

Further, since $\sum_{i=2}^m u_i^2 \leq \eta^2$, and in particular $u_i^2 \leq \eta^2$ for each $i \in \{2, \ldots, m\}$, we get that $\sum_{i=2}^m u_i^4 \leq \eta^2 \sum_{i=2}^m u_i^2 \leq \eta^4$. Continuing from equation (14), we have that

$$\tilde{F}_g(Z_1) - \tilde{F}_g(u) \geq \frac{1}{4}\alpha_{\min}\beta_{\min}^4 c_{\min}\eta^2 - \frac{1}{2}\alpha_{\max}\beta_{\max}^4 c_{\max}\eta^4 - 2\epsilon\ .$$

For any $\eta \in \left(\sqrt{\frac{16\epsilon}{\alpha_{\min}\beta_{\min}^4 c_{\min}}}, \sqrt{\frac{\alpha_{\min}\beta_{\min}^4 c_{\min}}{4\alpha_{\max}\beta_{\max}^4 c_{\max}}}\right)$, we obtain that

$$\tilde{F}_g(Z_1) - \tilde{F}_g(u) \geq \frac{1}{4}\alpha_{\min}\beta_{\min}^2 c_{\min}\eta^2 - \frac{1}{8}\alpha_{\min}\beta_{\min}^2 c_{\min}\eta^2 - 2\epsilon > 0\ .$$

Let $N = \{w \in S^{m-1} : \|Z_1 - w\| < \eta\}$. We note that for each $u \in \partial N$ (the boundary of N), $\tilde{F}_g(Z_1) > \tilde{F}_g(u)$. In particular, $v = \arg\max\{\tilde{F}_g(u) : u \in N \cup \partial N\}$ is contained in the interior of N. In particular, this choice of v is a local maximum of $\tilde{F}_g$. Using the lower bound on the arbitrary choice of η, we see that there exists v a local maximum of $\tilde{F}_g$ such that $\|Z_1 - v\| \leq \sqrt{\frac{16\epsilon}{\alpha_{\min}\beta_{\min}^4 c_{\min}}}$. $\square$

We now proceed with demonstrating that a small perturbation $\tilde{F}_g$ of F_g has no spurious maxima. For this, we enforce that the perturbation of the derivatives of $\tilde{F}_g$ is small so that we may ensure that the second derivative for extrema always implies that for u sufficiently far from the directions $\{\pm Z_i\}_{i=1}^m$, then $\tilde{F}(u)$ is not a maximum.

Lemma 30. *Suppose that F_g is $c_{\min}$-$c_{\max}$ robust. Then, for each $i \in [m]$ we have that*

$$|v_i| \cdot |2h_i'(v_i^2) - \langle \nabla F_g(v), v\rangle| \leq \|\nabla(F_g \circ \pi_v^{-1})(0)\| \ .$$

Proof. As an implication of Fact 27, we have that

$$\|\nabla(F_g \circ \pi_v^{-1})(\pi_v(v))\| = \|\nabla F_g(v) - \langle \nabla F_g(v), v\rangle v\| = \left\| \sum_{i=1}^m [2h_i'(v_i^2) - \langle \nabla F_g(v), v\rangle] v_i Z_i \right\| \ .$$

It follows that for each $i \in [m]$, $|(2h_i'(v_i^2) - \langle \nabla F_g(v), v\rangle)v_i| = |v_i| \cdot |2h_i'(v_i^2) - \langle \nabla F_g(v), v\rangle| \leq \|\nabla(F_g \circ \pi_v^{-1})(0)\|$. □

Lemma 31. *Suppose that F_g is $c_{\min}$-$c_{\max}$ robust, that $\eta > 0$, and that $v \in S^{m-1}$ has indices $i_1 \neq i_2$ such that $|v_{i_1}| > \eta$, $|v_{i_2}| > \eta$. If $\|\nabla(F_g \circ \pi_v^{-1})(0)\| < 2\alpha_{\min}\beta_{\min}^4 c_{\min}\eta^3$, then there exists a unit vector $x \in \mathcal{R}(\pi_v)$ such that $x^T \mathcal{H}(F_g \circ \pi_v^{-1})(0)x > 2\alpha_{\min}\beta_{\min}^4 c_{\min}\eta^2$.*

Proof. We let $x \in \mathcal{R}(\pi_v)$ be a unit vector and set $\xi = \mathrm{rep}(x) - v$ (which is also a unit vector). Using Fact 27 we obtain:

$$\begin{aligned} x^T[\mathcal{H}(F_g \circ \pi_v^{-1})(0)]x &= \xi^T[\mathcal{H}F_g(v) - \langle \nabla F_g(v), v\rangle I]\xi \\ &= \xi^T \left[\sum_{i=1}^m [4h_i''(u_i^2)u_i^2 + 2h_i'(u_i^2) - \langle \nabla F_g(v), v\rangle] Z_i Z_i^T \right] \xi \ . \end{aligned}$$

We may choose x such that $\xi \in \mathrm{span}(Z_{i_1}, Z_{i_2})$ is a unit vector. Then we get:

$$\begin{aligned} x^T[\mathcal{H}(F_g \circ \pi_v^{-1})(0)]x &= \xi^T \left[\sum_{i \in \{i_1, i_2\}} [4h_i''(v_i^2)v_i^2 + 2h_i'(v_i^2) - \langle \nabla F_g(v), v\rangle] Z_i Z_i^T \right] \xi \\ &\geq \xi^T \left[\sum_{i \in \{i_1, i_2\}} [4\alpha_{\min}\beta_{\min}^4 c_{\min}\eta^2 - \frac{\|\nabla(F_g \circ \pi_v^{-1})(0)\|}{\eta}] Z_i Z_i^T \right] \xi \\ &= 4\alpha_{\min}\beta_{\min}^4 c_{\min}\eta^2 - \frac{\|\nabla(F_g \circ \pi_v^{-1})(0)\|}{\eta} > 2\alpha_{\min}\beta_{\min}^4 c_{\min}\eta^2 \ . \end{aligned}$$

In the above, the first inequality uses Lemma 30 to obtain the second summand and a bound on $4h_i''(v_i^2)v_i^2$ (based on $|v_i| > \eta$ and equation (13)) to obtain the first summand. □

Proposition 32. *Suppose that F_g is $c_{\min}$-$c_{\max}$ robust, and that $\tilde{F}_g$ is a perturbation of F_g such that for some strictly positive constants ϵ_1 and ϵ_2, we have the following uniform bounds over all $u \in S^{m-1}$: $\|\nabla(\tilde{F}_g \circ \pi_u^{-1})(0) - \nabla(F_g \circ \pi_u^{-1})(0)\| < \epsilon_1$ and $\|\mathcal{H}(\tilde{F}_g \circ \pi_u^{-1})(0) - \mathcal{H}(F_g \circ \pi_u^{-1})(0)\| < \epsilon_2$. Then, for every v a local maximum of $\tilde{F}_g$ with respect to S^{m-1}, there exists $i \in [m]$ such that*

$$\|\mathrm{sign}(v_i)Z_i - v\| \leq \max\left(\left(\frac{\sqrt{2}\epsilon_1}{\alpha_{\min}\beta_{\min}^4 c_{\min}} \right)^{\frac{1}{3}} \sqrt{m}, \sqrt{\frac{\epsilon_2 m}{\alpha_{\min}\beta_{\min}^4 c_{\min}}} \right) \ .$$

Proof. We first fix a $v \in S^{m-1}$ and define $\Delta(v) := \min_{i \in [m]} \|\mathrm{sign}(v_i)Z_i - v\|$. Our first aim is to demonstrate that when $\Delta(v)$ is well separated from 0, there exist distinct $i_1, i_2 \in [m]$ such that $|v_{i_1}|$ and $|v_{i_2}|$ are also well separated from 0 so that we may apply Lemma 31.

Fixing a $j \in [m]$, we see that

$$\Delta(v)^2 \leq \|\text{sign}(v_j)Z_j - v\|^2 = \sum_{i \neq j} v_i^2 + (1 - |v_j|)^2 = \sum_{i=1}^m v_i^2 + 1 - 2|v_j| = 2 - 2|v_j| \ .$$

Rearranging terms, we obtain that $1 - |v_j| \geq \frac{1}{2}\Delta(v)^2$. Using again that $\sum_{i=1}^m v_i^2 = 1$, we see that

$$\sum_{i \neq j} v_i^2 = 1 - v_j^2 \geq (1 + |v_j|)(1 - |v_j)| \geq \frac{1}{2}\Delta(v)^2 \ .$$

It follows that there exists $i_1 \in [m] \setminus \{j\}$ such that $|v_{i_1}| \geq \frac{\Delta(v)}{\sqrt{2m}}$. Further, repeating the above construction with the choice of $j = i_1$ implies the existence of $i_2 \neq i_1$ such that $|v_{i_2}| \geq \frac{\Delta(v)}{\sqrt{2m}}$.

We proceed in arguing by contradiction. Suppose for the sake of contradiction that $v \in S^{m-1}$ is such that

$$\Delta(v) > \max\left(\left(\frac{\sqrt{2}\epsilon_1}{\alpha_{\min}\beta_{\min}^4 c_{\min}}\right)^{\frac{1}{3}}\sqrt{m}, \sqrt{\frac{\epsilon_2 m}{\alpha_{\min}\beta_{\min}^4 c_{\min}}}\right)$$

and that v meets the first order conditions for an extrema of $\tilde{F}_g$ on S^{m-1}. In particular, $\nabla(\tilde{F}_g \circ \pi_v^{-1})(0) = 0$. We set $\eta = \frac{\Delta(v)}{\sqrt{2m}}$, noting that there exists $i_1, i_2 \in [m]$ distinct such that $v_{i_1} > \eta$ and $v_{i_2} > \eta$. As $\epsilon_1 < \frac{1}{m\sqrt{2m}}\Delta(v)^3\alpha_{\min}\beta_{\min}^4 c_{\min} = 2\alpha_{\min}\beta_{\min}^4 c_{\min}\eta^3$, we obtain that $\|\nabla(F_g \circ \pi_v^{-1})(0)\| \leq \|\nabla(\tilde{F}_g \circ \pi_v^{-1})(0)\| + \epsilon_1 < 2\alpha_{\min}\beta_{\min}^4 c_{\min}\eta^3$.

By Lemma 31, there exists $x \in \mathcal{R}(\pi_v(0))$ that $x^T[\mathcal{H}(F_g \circ \pi_v^{-1})(0)]x > 2\alpha_{\min}\beta_{\min}^4 c_{\min}\eta^2 > \alpha_{\min}\beta_{\min}^4 c_{\min} \cdot \frac{\Delta(v)^2}{m}$. From the perturbation bounds, we obtain that

$$x^T[\mathcal{H}(\tilde{F}_g \circ \pi_v^{-1})(0)]x \geq x^T[\mathcal{H}(F_g \circ \pi_v^{-1})(0)]x - \epsilon_2 > \alpha_{\min}\beta_{\min}^4 c_{\min} \cdot \frac{\Delta(v)^2}{m} - \epsilon_2 > 0 \ .$$

By the second derivative conditions of extrema, this contradicts that v is a local maximum of $\tilde{F}_g \circ \pi_v^{-1}$ (and hence of $\tilde{F}_g$ with respect to S^{m-1}). It follows that for v a local maximum of S^{m-1}, we have $\Delta(v) \leq \max\left(\left(\frac{\sqrt{2}\epsilon_1}{\alpha_{\min}\beta_{\min}^4 c_{\min}}\right)^{\frac{1}{3}}\sqrt{m}, \sqrt{\frac{\epsilon_2 m}{\alpha_{\min}\beta_{\min}^4 c_{\min}}}\right)$. □

Our main results in this subsection are to demonstrate the existence local maxima of a perturbation of F_g near each of the directions $\pm Z_i$, and to demonstrate that there are no spurious maxima. More formally, we have the following theorem.

Theorem 33. *Suppose that F_g is $c_{\min}$-$c_{\max}$ robust, and that $\tilde{F}_g$ is a perturbation of F_g such that for some strictly positive constants ϵ_0, ϵ_1, and ϵ_2, we have the following uniform bounds over all $u \in S^{m-1}$: $|\tilde{F}_g(u) - F_g(u)| < \epsilon_0$, $\|\nabla(\tilde{F}_g \circ \pi_u^{-1})(0) - \nabla(F_g \circ \pi_u^{-1})(0)\| < \epsilon_1$, and $\|\mathcal{H}(\tilde{F}_g \circ \pi_u^{-1})(0) - \mathcal{H}(F_g \circ \pi_u^{-1})(0)\| < \epsilon_2$. Suppose that $\epsilon_0 < \frac{\alpha_{\min}^2\beta_{\min}^8 c_{\min}^2}{64\alpha_{\max}\beta_{\max}^4 c_{\max}}$. Then $\tilde{F}_g$ satisfies the following existence and localization guarantees for its maxima.*

- *(Existence) For every $i \in [m]$ and each sign $s \in \{\pm 1\}$, there exists $v \in S^{m-1}$ a local maximum of F_g with respect to S^{m-1} such that $\|sZ_i - v\| \leq \max\left(\left(\frac{\sqrt{2}\epsilon_1}{\alpha_{\min}\beta_{\min}^4 c_{\min}}\right)^{\frac{1}{3}}\sqrt{m}, \sqrt{\frac{\epsilon_2 m}{\alpha_{\min}\beta_{\min}^4 c_{\min}}}\right)$.*
- *(No Spurious Maxima) If v is a local maximum of $\tilde{F}_g$ with respect to S^{m-1}, then there exists $i \in [m]$ and a sign value $s \in \{\pm 1\}$ such that $\|sZ_i - v\| \leq \max\left(\left(\frac{\sqrt{2}\epsilon_1}{\alpha_{\min}\beta_{\min}^4 c_{\min}}\right)^{\frac{1}{3}}\sqrt{m}, \sqrt{\frac{\epsilon_2 m}{\alpha_{\min}\beta_{\min}^4 c_{\min}}}\right)$.*

Proof. The no spurious maxima result is a restatement of Proposition 32.

For the existence result, we apply Lemma 28 to obtain that for any choice of $i \in [m]$ and $s \in \{\pm 1\}$, there exists a local maximum v of $\tilde{F}_g$ on the sphere such that $\|sZ_i - v\| \leq \sqrt{\frac{16\epsilon_0}{\alpha_{\min}\beta_{\min}^4 c_{\min}}} < \frac{1}{2}$. In particular, $\langle v, sZ_i\rangle = \sqrt{1 - \|P_{Z_i^\perp} v\|^2} > \sqrt{1 - (\frac{1}{2})^2} \geq \frac{\sqrt{3}}{2}$. For any Z_j with $j \neq i$, we have that $|\langle v, Z_j\rangle| \leq \|P_{Z_i^\perp} v\| \leq \frac{1}{2}$. By contraposition on the implication ($\|sZ_i - v\| < \frac{1}{2}$ implies $\langle v, sZ_i\rangle > \frac{\sqrt{3}}{2}$), it follows that for any choice of $j \neq i$ and $s' \in \{\pm 1\}$, $\|s'Z_j - v\| > \frac{1}{2}$. In this sense, v is uniquely associated with sZ_i.

When $\max\left(\left(\frac{\sqrt{2}\epsilon_1}{\alpha_{\min}\beta_{\min}^4 c_{\min}}\right)^{\frac{1}{3}}\sqrt{m}, \sqrt{\frac{\epsilon_2 m}{\alpha_{\min}\beta_{\min}^4 c_{\min}}}\right) \geq \frac{1}{2}$, then the existence result if trivial. When $\max\left(\left(\frac{\sqrt{2}\epsilon_1}{\alpha_{\min}\beta_{\min}^4 c_{\min}}\right)^{\frac{1}{3}}\sqrt{m}, \sqrt{\frac{\epsilon_2 m}{\alpha_{\min}\beta_{\min}^4 c_{\min}}}\right) < \frac{1}{2}$, then the choice of v a local maximum of F_g such that $\|sZ_i - v\| < \frac{1}{2}$ must satisfy $\|sZ_i - v\| \leq \max\left(\left(\frac{\sqrt{2}\epsilon_1}{\alpha_{\min}\beta_{\min}^4 c_{\min}}\right)^{\frac{1}{3}}\sqrt{m}, \sqrt{\frac{\epsilon_2 m}{\alpha_{\min}\beta_{\min}^4 c_{\min}}}\right)$ using Proposition 32 and noting that $\|s'Z_j - v\| \leq \max\left(\left(\frac{\sqrt{2}\epsilon_1}{\alpha_{\min}\beta_{\min}^4 c_{\min}}\right)^{\frac{1}{3}}\sqrt{m}, \sqrt{\frac{\epsilon_2 m}{\alpha_{\min}\beta_{\min}^4 c_{\min}}}\right) < \frac{1}{2}$ cannot hold for any other choice of $s' \in \{\pm 1\}$ and $j \neq i$. □

E.3 Contrast perturbation under the spectral embedding

In this subsection, we combine the bounds from Appendix D with our perturbation results for F_g in order to demonstrate that within our spectral clustering application, the maxima structure of F_g is robust to a perturbation of the graph Laplacian. We assume throughout this subsection that g is $(c_{\min}, c_{\max}, D)$-robust. For simplicity, we work with the standard graph Laplacian[3] L. We further assume that $\tilde{L} = L + H$ is a perturbation of L (H is viewed as being small). Under Theorem 22, we let $\tilde{X}$ provide the embedded data from $\tilde{L}$, and we let X contain the embedded data from L using a choice of the basis for $\mathcal{N}(L)$ such that $\frac{1}{\sqrt{n}}\|X - \tilde{X}\| \leq \frac{2\|H\|}{\delta(L) - \|H\|}$ holds.

From these embeddings, we define F_g and $\tilde{F}_g$ according to equation (3). In particular, $F_g(u) = \frac{1}{n}\sum_{i=1}^n g(|\langle u, x_{i\cdot}\rangle|)$ and $\tilde{F}_g(u) = \frac{1}{n}\sum_{i=1}^n g(|\langle u, \tilde{x}_{i\cdot}\rangle|)$. We also recall from equation (4) that F_g may be equivalently written as $F_g(u) = \sum_{j=1}^m w_j g(\frac{1}{\sqrt{w_j}}|\langle u, Z_j\rangle|)$ for some hidden orthonormal basis $Z_1, \ldots, Z_m$. In particular, F_g has associated constants $\alpha_j = w_j$ and $\beta_j = \frac{1}{\sqrt{w_j}}$, and $\tilde{F}_g$ is its perturbation.

We now bound the perturbation error in $\tilde{F}_g$ arising from the embedding error in $\tilde{X}$.

Lemma 34. *Let $E > 0$ be a fixed constant. There exists a constant $C > 0$ depending only on E such that the following holds. If $\|H\| \leq \delta(L)\min(\frac{1}{2}, \frac{E}{\sqrt{n}})$, then for any $u \in S^{m-1}$ and unit vector $v \in T_u S^{m-1}$, we have:*

1. $|\tilde{F}_g(u) - F_g(u)| \leq C\frac{c_{\max}}{w_{\min}^{3/2}} \cdot \frac{\|H\|}{\delta(L)}$.
2. $|\langle \nabla(\tilde{F}_g \circ \pi_u^{-1})(0) - \nabla(F_g \circ \pi_u^{-1})(0), v\rangle| \leq C\frac{c_{\max}}{w_{\min}^{3/2}} \cdot \frac{\|H\|}{\delta(L)}$.
3. $|v^T[\mathcal{H}(\tilde{F}_g \circ \pi_u^{-1})(0) - \mathcal{H}(F_g \circ \pi_u^{-1})(0)]v| \leq C[\frac{c_{\max}}{w_{\min}^{3/2}} + \frac{D}{w_{\min}^{5/2}}] \cdot \frac{\|H\|}{\delta(L)}$.

[3]We exclude the symmetric normalized Laplacian L_{sym} in this section in order to simplify the analysis. The ray structure which arises from the spectral embedding with L_{sym} makes it so that the distance of points from the origin depends not only on the cluster weights $w_1, \ldots, w_m$, but also one the degrees of each individual graph vertex in accordance with Lemma 15.

Proof. Throughout, we will denote by $h : \mathbb{R} \to \mathbb{R}$ the function $h(t) = g(\sqrt{|t|})$. As such, we may write $F_g(u) = \frac{1}{n}\sum_{i=1}^n h(\langle u, x_{i\cdot}\rangle^2)$ and $\tilde{F}_g(u) = \frac{1}{n}\sum_{i=1}^n h(\langle u, \tilde{x}_{i\cdot}\rangle^2)$. By assumption R1, $h''(t) \in [c_{\min}, c_{\max}]$ for all $t \in \mathbb{R}$. In order to compress notation, we will denote the expressions $w := \frac{1}{\sqrt{w_{\min}}} = \max_{i\in[n]}\|x_{i\cdot}\|$ (by Proposition 4) and $\epsilon := \frac{2\|H\|}{\delta(L)-\|H\|}$ is the bound from Theorem 22. Notice that by our assumptions,

$$\epsilon \leq 4\frac{\|H\|}{\delta(L)} \leq \min(2, \frac{E}{\sqrt{n}}) \ . \tag{15}$$

Proof of part 1. We now bound $|\tilde{F}_g(u) - F_g(u)|$.

$$\begin{aligned}
|\tilde{F}_g(u) - F_g(u)| &= \frac{1}{n}\left|\sum_{i=1}^n [h(\langle u, \tilde{x}_{i\cdot}\rangle^2) - h(\langle u, x_{i\cdot}\rangle^2)]\right| \\
&= \frac{1}{n}\left|\sum_{i=1}^n [h'(\langle u, x_{i\cdot}\rangle^2)(\langle u, \tilde{x}_{i\cdot}\rangle^2 - \langle u, x_{i\cdot}\rangle^2) + \frac{1}{2}h''(c_i)(\langle u, \tilde{x}_{i\cdot}\rangle^2 - \langle u, x_{i\cdot}\rangle^2)^2]\right| \\
&\leq \frac{1}{n}\sum_{i=1}^n \left[|c_{\max}w^2(\langle u, \tilde{x}_{i\cdot}\rangle^2 - \langle u, x_{i\cdot}\rangle^2)| + |\frac{1}{2}c_{\max}(\langle u, \tilde{x}_{i\cdot}\rangle^2 - \langle u, x_{i\cdot}\rangle^2)^2|\right] \ .
\end{aligned} \tag{16}$$

In the above, each $c_i \in [\langle u, x_{i\cdot}\rangle^2, \langle u, \tilde{x}_{i\cdot}\rangle^2]$ comes from the error term in the Taylor expansion of h. The final inequality uses the triangle inequality, assumption R1, and $w \geq |\langle u, x_{i\cdot}\rangle|$ for all $i \in [m]$.

Noting that each $\langle u, x_{i\cdot}\rangle \leq \|x_{i\cdot}\| \leq \frac{1}{\sqrt{w_{\min}}} = w$, writing $\langle u, \tilde{x}_{i\cdot}\rangle = (\langle u, \tilde{x}_{i\cdot}\rangle - \langle u, x_{i\cdot}\rangle) + \langle u, x_{i\cdot}\rangle$, and using the ϵ-bound from Theorem 22, we obtain a worst case bound of

$$|\langle u, \tilde{x}_{i\cdot}\rangle| \leq w + \epsilon\sqrt{n} \ . \tag{17}$$

Further, it can be seen that

$$\frac{1}{n}\sum_{i=1}^n |\langle u, \tilde{x}_{i\cdot} - x_{i\cdot}\rangle| \leq \frac{1}{n}\|\tilde{X} - X\|_{2,1} \leq \frac{1}{\sqrt{n}}\|\tilde{X} - X\| \leq \epsilon \ . \tag{18}$$

where $\|\cdot\|_{2,1}$ is the induced norm defined by $\|A\|_{2,1} := \max\{\|Ay\|_1 : \|y\|_2 = 1\}$. The second inequality uses that for any vector $y \in \mathbb{R}^n$, $\|y\|_1 \leq \|y\|_2\sqrt{n}$. The final inequality uses the bound from Theorem 22. Also,

$$\frac{1}{n}\sum_{i=1}^n |\langle u, \tilde{x}_{i\cdot}\rangle^2 - \langle u, x_{i\cdot}\rangle^2| = \frac{1}{n}\sum_{i=1}^n |\langle u, \tilde{x}_{i\cdot} - x_{i\cdot}\rangle\langle u, \tilde{x}_{i\cdot} + x_{i\cdot}\rangle| \leq (2w + \epsilon\sqrt{n})\epsilon \tag{19}$$

by using equations (17) and (18), and by recalling that $w = \frac{1}{\sqrt{w_{\min}}}$ is an upper bound for $\|x_{i\cdot}\|$. However, for a slight variation, we get a nicer bound:

$$\begin{aligned}
\frac{1}{n}\sum_{i=1}^n (\langle u, \tilde{x}_{i\cdot}\rangle^2 - \langle u, x_{i\cdot}\rangle^2)^2 &= \frac{1}{n}\sum_{i=1}^n \langle u, \tilde{x}_{i\cdot} - x_{i\cdot}\rangle^2\langle u, \tilde{x}_{i\cdot} + x_{i\cdot}\rangle^2 \\
&\leq \frac{1}{n}(2w + \epsilon\sqrt{n})^2\|\tilde{X} - X\|^2 \leq (2w + \epsilon\sqrt{n})^2\epsilon^2
\end{aligned} \tag{20}$$

Continuing from equation (16) and applying equation (19) and (20), we obtain

$$|\tilde{F}_g(u) - F_g(u)| \leq c_{\max}w^2(2w + \epsilon\sqrt{n})\epsilon + \frac{1}{2}c_{\max}(2w + \epsilon\sqrt{n})^2\epsilon^2 \ .$$

Using equation (15), $|\tilde{F}_g(u) - F_g(u)|$ is upper bounded by $O(w^3 c_{\max}\|H\|/\delta(L))$. It will be convenient later to have a specific constant here (which we make no attempt to optimize). Since $\epsilon\sqrt{n} \le E$, we obtain:

$$\begin{aligned}
|\tilde{F}_g(u) - F_g(u)| &\le c_{\max}[w^2(2w+E) + \frac{1}{2}(4w^2 + 4wE + E^2)\epsilon]\epsilon \\
&\le c_{\max}[w^2(2w+E) + \frac{1}{2}(4w^2 + 4wE + E^2)E]\epsilon \\
&\le c_{\max}[2w^3 + 3Ew^2 + 2E^2 w + \frac{1}{2}E^3]\epsilon \\
&\le 8c_{\max}\max(w,E)^3\epsilon \le 8c_{\max}(1+E)^3 w^3\epsilon \\
&\le 32(1+E)^3 \frac{c_{\max}\|H\|}{w_{\min}^{3/2}\delta(L)} \qquad (21)
\end{aligned}$$

Note that in the last inequality, we use equation (15).

Proof of part 2. We now bound $|\langle \nabla(\tilde{F}_g \circ \pi_u^{-1})(0) - \nabla(\tilde{F}_g \circ \pi_u^{-1})(0), v\rangle|$ for some choice of $v \in T_u S^{m-1}$. We set $\xi = \mathrm{rep}(v) - u$ another unit vector, and we obtain:

$$\begin{aligned}
&|\langle \nabla(\tilde{F}_g \circ \pi_u^{-1})(0) - \nabla(F_g \circ \pi_u^{-1})(0), v\rangle| = |\langle \nabla\tilde{F}_g(u) - \nabla F_g(u), \xi\rangle| \\
&= \frac{2}{n}\left|\left\langle \sum_{i=1}^n [h'(\langle u, \tilde{x}_{i\cdot}\rangle^2)\langle u, \tilde{x}_{i\cdot}\rangle \tilde{x}_{i\cdot}^T - h'(\langle u, x_{i\cdot}\rangle^2)\langle u, x_{i\cdot}\rangle x_{i\cdot}],\ \xi \right\rangle\right| \\
&\le \frac{2}{n}\sum_{i=1}^n \Big| h'(\langle u, \tilde{x}_{i\cdot}\rangle^2)\langle u, \tilde{x}_{i\cdot}\rangle[\langle \xi, \tilde{x}_{i\cdot}\rangle - \langle \xi, x_{i\cdot}\rangle] + h'(\langle u, \tilde{x}_{i\cdot}\rangle^2)[\langle u, \tilde{x}_{i\cdot}\rangle - \langle u, x_{i\cdot}\rangle]\langle \xi, x_{i\cdot}\rangle \\
&\qquad + [h'(\langle u, \tilde{x}_{i\cdot}\rangle^2) - h'(\langle u, x_{i\cdot}\rangle^2)]\langle u, x_{i\cdot}\rangle\langle \xi, x_{i\cdot}\rangle \Big| \\
&\le \frac{2}{n}\sum_{i=1}^n \left|c_{\max}\langle u, \tilde{x}_{i\cdot}\rangle^2\langle u, \tilde{x}_{i\cdot}\rangle\langle \xi, \tilde{x}_{i\cdot} - x_{i\cdot}\rangle\right| + \left|c_{\max}\langle u, \tilde{x}_{i\cdot}\rangle^2\langle u, \tilde{x}_{i\cdot} - x_{i\cdot}\rangle\langle \xi, x_{i\cdot}\rangle\right| \\
&\qquad + \left|c_{\max}[\langle u, \tilde{x}_{i\cdot}\rangle^2 - \langle u, x_{i\cdot}\rangle^2]\langle u, x_{i\cdot}\rangle\langle \xi, x_{i\cdot}\rangle\right| \qquad (22)
\end{aligned}$$

where the second inequality uses the triangle inequality and assumption R1.

Using the bounds $w \ge \|x_{i\cdot}\|$, equation (17) and equation (18), we simplify (22) as:

$$\begin{aligned}
|\langle \nabla(\tilde{F}_g \circ \pi_u^{-1})(0) - \nabla(F_g \circ \pi_u^{-1})(0), v\rangle| &\le 2c_{\max}\left[(w+\epsilon\sqrt{n})^2[(w+\epsilon\sqrt{n})\epsilon + w\epsilon] + w^2(2w+\epsilon\sqrt{n})\epsilon\right] \\
&= 2c_{\max}[(w+\epsilon\sqrt{n})^2 + w^2][2w+\epsilon\sqrt{n}]\epsilon\ .
\end{aligned}$$

In particular, $|\langle \nabla(\tilde{F}_g \circ \pi_u^{-1})(0) - \nabla(F_g \circ \pi_u^{-1})(0), v\rangle|$ is upper bounded by $O(c_{\max} w^3\|H\|/\delta(L))$.

Proof of part 3. We set $\xi = \mathrm{rep}(v) - u$ another unit vector, and we obtain:

$$\begin{aligned}
&|v^T[\mathcal{H}(\tilde{F}_g \circ \pi_u^{-1})(0) - \mathcal{H}(F_g \circ \pi_u^{-1})(0)]v| \\
&= |\xi^T[\mathcal{H}\tilde{F}_g(u) - \langle \nabla F_g(u), u\rangle I - \mathcal{H}F_g(u) + \langle \nabla F_g(u), u\rangle I]\xi| \\
&= \frac{1}{n}\Big|\xi^T\Big[\sum_{i=1}^n \Big([4h''(\langle u, \tilde{x}_{i\cdot}\rangle^2)\langle u, \tilde{x}_{i\cdot}\rangle^2 + 2h'(\langle u, \tilde{x}_{i\cdot}\rangle^2)]\tilde{x}_{i\cdot}^T\tilde{x}_{i\cdot} - 2h'(\langle u, \tilde{x}_{i\cdot}\rangle^2)\langle u, \tilde{x}_{i\cdot}\rangle^2 I \\
&\qquad - [4h''(\langle u, x_{i\cdot}\rangle^2)\langle u, x_{i\cdot}\rangle^2 + 2h'(\langle u, x_{i\cdot}\rangle^2)]x_{i\cdot}^T x_{i\cdot} + 2h'(\langle u, x_{i\cdot}\rangle^2)\langle u, x_{i\cdot}\rangle^2 I\Big)\xi\Big|\ . \qquad (23)
\end{aligned}$$

We now expand several of the summands in equation (23) separately.

$$\begin{aligned}
&\frac{4}{n}\Big|\sum_{i=1}^{n}[h''(\langle u,\tilde{x}_{i\cdot}\rangle^2)\langle u,\tilde{x}_{i\cdot}\rangle^2\langle\xi,\tilde{x}_{i\cdot}\rangle^2 - h''(\langle u,x_{i\cdot}\rangle^2)\langle u,x_{i\cdot}\rangle^2\langle\xi,x_{i\cdot}\rangle^2]\Big| \\
&= \frac{4}{n}\Big|\sum_{i=1}^{n}\Big(h''(\langle u,\tilde{x}_{i\cdot}\rangle^2)\langle u,\tilde{x}_{i\cdot}\rangle^2[\langle\xi,\tilde{x}_{i\cdot}\rangle^2 - \langle\xi,x_{i\cdot}\rangle^2] + h''(\langle u,\tilde{x}_{i\cdot}\rangle^2)[\langle u,\tilde{x}_{i\cdot}\rangle^2 - \langle u,x_{i\cdot}\rangle^2]\langle\xi,x_{i\cdot}\rangle^2 \\
&\qquad + [h''(\langle u,\tilde{x}_{i\cdot}\rangle^2) - h''(\langle u,x_{i\cdot}\rangle^2)]\langle u,x_{i\cdot}\rangle^2\langle\xi,x_{i\cdot}\rangle^2\Big)\Big| \\
&\le \frac{4}{n}\sum_{i=1}^{n}\Big[c_{\max}(w+\epsilon\sqrt{n})^2|\langle\xi,\tilde{x}_{i\cdot}\rangle^2 - \langle\xi,x_{i\cdot}\rangle^2| + c_{\max}|\langle u,\tilde{x}_{i\cdot}\rangle^2 - \langle u,x_{i\cdot}\rangle^2|w^2 \\
&\qquad + D|\langle u,\tilde{x}_{i\cdot}\rangle^2 - \langle u,x_{i\cdot}\rangle^2|w^4\Big] \\
&\le 4c_{\max}[(w+\epsilon\sqrt{n})^2 + w^2 + Dw^4](2w+\epsilon\sqrt{n})\epsilon\ .
\end{aligned}$$

In the above, the first inequality uses the mean value theorem, equation (17), and the triangular inequality. The second inequality uses equation (19).

We now bound a second summand grouping from equation (23):

$$\begin{aligned}
&\frac{2}{n}\left|\sum_{i=1}^{n}[h'(\langle u,\tilde{x}_{i\cdot}\rangle^2)\langle\xi,\tilde{x}_{i\cdot}^2\rangle - h'(\langle u,x_{i\cdot}\rangle^2)\langle\xi,x_{i\cdot}^2\rangle]\right| \\
&= \frac{2}{n}\sum_{i=1}^{n}|h'(\langle u,\tilde{x}_{i\cdot}\rangle^2)[\langle\xi,\tilde{x}_{i\cdot}\rangle^2 - \langle\xi,x_{i\cdot}\rangle^2] + [h'(\langle u,\tilde{x}_{i\cdot}\rangle^2) - h'(\langle u,x_{i\cdot}\rangle^2)]\langle\xi,x_{i\cdot}\rangle^2| \\
&\le \frac{2}{n}\sum_{i=1}^{n}\left[c_{\max}\langle u,\tilde{x}_{i\cdot}\rangle^2|\langle\xi,\tilde{x}_{i\cdot}\rangle^2 - \langle\xi,x_{i\cdot}\rangle^2| + c_{\max}|\langle u,\tilde{x}_{i\cdot}\rangle^2 - \langle u,x_{i\cdot}\rangle^2|\langle\xi,x_{i\cdot}\rangle^2\right] \\
&\le 2c_{\max}((w+\epsilon\sqrt{n})^2 + w^2)(2w+\epsilon\sqrt{n})\epsilon\ , \qquad (24)
\end{aligned}$$

where the first inequality uses the mean value theorem, and the second inequality uses equation (19), equation (17), and that $w \ge \|x_{i\cdot}\|$.

We c bound the final summand grouping from equation (23) as

$$\frac{2}{n}\left|\sum_{i=1}^{n}[h'(\langle u,x_{i\cdot}\rangle^2)\langle u,x_{i\cdot}\rangle^2 - h'(\langle u,\tilde{x}_{i\cdot}\rangle^2)\langle u,\tilde{x}_{i\cdot}\rangle^2]\right| \le 2c_{\max}((w+\epsilon\sqrt{n})^2 + w^2)(2w+\epsilon\sqrt{n})\epsilon$$

by using the same argument as in (24) but replacing ξ with u. Collecting all of these terms, we see that $|v^T[\mathcal{H}(\tilde{F}_g \circ \pi_u^{-1})(0) - \mathcal{H}(F_g \circ \pi_u^{-1})(0)]v|$ is upper bounded by $O(c_{\max}(w^3 + Dw^5)\epsilon) = O(c_{\max}(w^3 + Dw^5)\|H\|/\delta(L))$. □

We now state and prove our main perturbation result demonstrating that the maxima of F_g are robust to a small perturbation of the Graph Laplacian from the setting of a similarity graph consisting of m connected components.

Theorem 35. *Given $E > 0$ a constant, there exists $C > 0$ a constant such that the following holds. Suppose that $\tilde{X} = (\tilde{x}_{\cdot 1}, \ldots, \tilde{x}_{\cdot m})$ contains the lowest m eigenvectors of $\tilde{L}$ scaled such that each $\|\tilde{x}_{\cdot i}\| = \frac{1}{\sqrt{n}}$. Let $X = (x_{\cdot 1}, \ldots, x_{\cdot m})$ be a scaled orthogonal basis of $\mathcal{N}(L)$ such that each $\|x_{\cdot i}\| = \frac{1}{\sqrt{n}}$ and $\frac{1}{\sqrt{n}}\|X - \tilde{X}\| \le \frac{2\|H\|}{\delta(L) - \|H\|}$. Let g be a $(c_{\min}, c_{\max}, D)$-robust contrast, and construct F_g from $x_{1\cdot}, \ldots, x_{n\cdot}$ and $\tilde{F}_g$ from $\tilde{x}_{1\cdot}, \ldots, \tilde{x}_{n\cdot}$ according to equation (3). Let $Z_1, \ldots, Z_m$*

denote the distinct hidden basis directions from the cleanly embedded points $\frac{x_{i\cdot}^T}{\|x_{i\cdot}\|}$. *If* $\|H\| \leq \delta(L)\min(\frac{E}{\sqrt{n}}, \frac{w_{\min}^{5.5}c_{\min}^2}{2048w_{\max}^5 c_{\max}(1+E)^3}, \frac{1}{2})$, *then the following hold:*

1. *(Maxima Existence) For every* $i \in [m]$ *and every* $s \in \{\pm 1\}$, *there exists a local maximum* v *of* $\tilde{F}_g$ *with respect to* S^{m-1} *such that* $\|sZ_i - v\| \leq C\big(\frac{\|H\|}{\delta(L)}\big)^{1/3}\sqrt{\frac{[c_{\max}+D]m}{c_{\min}} \cdot \frac{w_{\max}^2}{w_{\min}^{7/2}}}$.

2. *(No Spurious Maxima) If* v *is a local maximum of* $\tilde{F}_g$ *with respect to* S^{m-1}, *then there exists* $i \in [m]$ *and* $s \in \{\pm 1\}$ *such that* $\|sZ_i - v\| \leq C\big(\frac{\|H\|}{\delta(L)}\big)^{1/3}\sqrt{\frac{[c_{\max}+D]m}{c_{\min}} \cdot \frac{w_{\max}^2}{w_{\min}^{7/2}}}$.

We note that the existence of X in Theorem 35 is guaranteed by Theorem 22. As Theorem 35 requires that $\|H\|$ be upper bounded as $O(1/\sqrt{n})$, it is worth noting that the eigenvalues and eigenprojections of kernel estimates concentrate with $O(1/\sqrt{n})$ error [18].

Proof. This is an exercise in collecting the bounds from Lemma 34 and Theorem 33. In particular, we define $\epsilon_0 := 32(1+E)^3\frac{c_{\max}\|H\|}{w_{\min}^{3/2}\delta(L)}$, $\epsilon_1 := C'\frac{c_{\max}\|H\|}{w_{\min}^{3/2}\delta(L)}$, and $\epsilon_2 := C'\Big[\frac{c_{\max}}{w_{\min}^{3/2}} + \frac{D}{w_{\min}^{5/2}}\Big] \cdot \frac{\|H\|}{\delta(L)}$ where C' is the constant from Lemma 34. Then Lemma 34 (combined with equation (21) for ϵ_0) implies that: (0) $|\tilde{F}_g(u) - F_g(u)| \leq \epsilon_0$ for all $u \in S^{n-1}$, (1) $|\langle \nabla\tilde{F}_g \circ \pi_u^{-1}(0) - \nabla F_g \circ \pi_u^{-1}(0), v\rangle| < \epsilon_1$ for all $u \in S^{m-1}$ and all unit vectors $v \in T_uS^{m-1}$, and (2) $|v^T[\mathcal{H}(\tilde{F}_g \circ \pi_u^{-1})(0) - \mathcal{H}(F_g \circ \pi_u^{-1})(0)]v| < \epsilon_2$ for all $u \in S^{m-1}$ and $v \in T_uS^{n-1}$.

Since $\|H\| \leq \frac{w_{\min}^{5.5}c_{\min}^2\delta(L)}{2048w_{\max}^5 c_{\max}(1+E)^3}$, it follows that $\epsilon_0 \leq \frac{w_{\min}^4 c_{\min}^2}{64w_{\max}^5 c_{\max}}$. In particular, we may apply Theorem 33 with $\alpha_{\min} = w_{\min}$, $\alpha_{\max} = w_{\max}$, $\beta_{\min} = w_{\max}^{-1/2}$, and $\beta_{\max} = w_{\min}^{-1/2}$ chosen to match the construction of F_g in the spectral embedding case (see equation (4)). Doing so, we obtain that for all $v \in S^{m-1}$ a maximum of $\tilde{F}_g$, there exists a choice of $s \in \{\pm 1\}$ and $i \in [dm]$ such that

$$\begin{aligned}
\|sZ_i - v\| &\leq \max\left(\left(\frac{\sqrt{2}\epsilon_1}{\alpha_{\min}\beta_{\min}^4 c_{\min}}\right)^{\frac{1}{3}}\sqrt{m}, \sqrt{\frac{\epsilon_2 m}{\alpha_{\min}\beta_{\min}^4 c_{\min}}}\right) \\
&\leq \max\left(\left(\frac{C'\sqrt{2}\frac{c_{\max}\|H\|}{w_{\min}^{3/2}\delta(L)}}{\frac{w_{\min}}{w_{\max}^2}c_{\min}}\right)^{\frac{1}{3}}\sqrt{m}, \sqrt{\frac{C'\Big[\frac{c_{\max}}{w_{\min}^{3/2}} + \frac{D}{w_{\min}^{5/2}}\Big] \cdot \frac{\|H\|}{\delta(L)}m}{\frac{w_{\min}}{w_{\max}^2}c_{\min}}}\right) \\
&= O\left(\Big(\frac{\|H\|}{\delta(L)}\Big)^{1/3}\sqrt{\frac{[c_{\max}+D]mw_{\max}^2}{c_{\min}w_{\min}^{7/2}}}\right). \qquad \square
\end{aligned}$$

F Derivatives on the Sphere

In this section, we compute the derivative formulas involving local coordinate system π_v which was introduced in section E. We now derive the needed first derivative formulas.

Lemma 36. *Let* $f \in \mathcal{C}^2(\mathbb{R}^n)$ *and let* $v \in S^{n-1}$. *Letting* $u(x) = \pi_v^{-1}(x)$ *denote the inverse change*

of coordinates, then the first two derivatives of $f \circ \pi_v^{-1}$ are given by

$$\partial_i(f \circ \pi_v^{-1})(x) = \sum_{k=1}^{n-1} \partial_{p_k} f(u) \left[\frac{\delta_{ik}}{\|(1,\ x^T)\|} - \frac{x_i x_k}{\|(1,\ x^T)\|^3} \right] - \frac{\partial_v f(u) x_i}{\|(1,\ x^T)\|^3} \tag{25}$$

$$\begin{aligned}
\partial_{ji}(f \circ \pi_v^{-1})(x) = &\sum_{\ell=1}^{n-1} \sum_{k=1}^{n-1} \partial_{p_\ell p_k} f(u) \left[\frac{\delta_{j\ell}}{\|(1,\ x^T)\|} - \frac{x_j x_\ell}{\|(1,\ x^T)\|^3} \right] \left[\frac{\delta_{ik}}{\|(1,\ x^T)\|} - \frac{x_i x_k}{\|(1,\ x^T)\|^3} \right] \\
&- \sum_{k=1}^{n-1} \frac{\partial_{v p_k} f(u)}{\|(1, x^T)\|^4} \left[\delta_{ik} - \frac{x_i x_k}{\|(1,\ x^T)\|^2} \right] x_j - \sum_{k=1}^{n-1} \frac{\partial_{p_k v} f(u)}{\|(1, x^T)\|^4} \left[\delta_{jk} - \frac{x_j x_k}{\|(1,\ x^T)\|^2} \right] x_i \\
&+ \frac{\partial_{vv} f(u) x_i x_j}{\|(1,\ x^T)\|^6} + \sum_{k=1}^{n-1} \partial_{p_k} f(u) \left[\frac{3 x_i x_j x_k}{\|(1, x^T)\|^5} - \frac{\delta_{kj} x_i + \delta_{ki} x_j + \delta_{ij} x_k}{\|(1, x^T)\|^3} \right] \\
&+ \partial_v f(u) \left[\frac{3 x_i x_j}{\|(1,\ x^T)\|^5} - \frac{\delta_{ij}}{\|(1,\ x^T)\|^3} \right]
\end{aligned} \tag{26}$$

Proof. We may write $f \circ \pi_v^{-1}$ simply as $f(u)$. By the chain rule, we get that $\partial_i(f \circ \pi_v^{-1}) = \sum_{k=1}^n \partial_{p_k} f(u) \langle \partial_i u, p_k \rangle$ where

$$\partial_i u(x) = -\frac{x_i}{\|(1,\ x^T)\|^3} v - \sum_{k=1}^{n-1} \frac{x_i x_k}{\|(1,\ x^T)\|^3} p_k + \frac{1}{\|(1,\ x^T)\|} p_i \ .$$

Combining the mentioned formulas gives equation (25). Continuing, we obtain the second derivative formula $\partial_{ji}(f \circ \pi_v^{-1}) = \sum_{\ell=1}^n \sum_{k=1}^n \partial_{p_\ell p_k} f(u) \langle \partial_j u, p_\ell \rangle \langle \partial_i u, p_k \rangle + \sum_{k=1}^n \partial_{p_k} f(u) \langle \partial_{ji} u, p_k \rangle$. But we have:

$$\partial_{ji} u(x) = \frac{3 x_i x_j [v + \sum_{k=1}^{n-1} x_k p_k]}{\|(1,\ x^T)\|^5} - \frac{\delta_{ij}(v + \sum_{k=1}^{n-1} x_k p_k) + x_i p_j + x_j p_i}{\|(1,\ x^T)\|^3} \ .$$

Putting the formulas together gives equation (26). □

Corollary 37. *Let $f \in \mathcal{C}^2(\mathbb{R}^n)$ and let $v \in S^{n-1}$. Letting $u(x) = \pi_v^{-1}(x)$ denote the inverse change of coordinates, then the first two derivatives of $f \circ \pi_v^{-1}$ evaluated at $\pi_v(v) = 0$ are*

$$\partial_i(f \circ \pi_v^{-1})(\pi_v(v)) = \partial_{p_i} f(v) \tag{27}$$

$$\partial_{ji}(f \circ \pi_v^{-1})(\pi_v(v)) = \partial_{p_j p_i} f(v) - \partial_v f(v) \delta_{ij} \tag{28}$$

We now derive the needed second derivative formulas.

Lemma 38. *Suppose that $x, y \in \mathcal{R}(\pi_v)$, and let $\xi = \mathrm{rep}(x) - v$ and $\zeta = \mathrm{rep}(y) - v$ be the corresponding translations in $T_v S^{n-1}$. Then, given a vector $u \in S_v^+$, we have the following operator equivalences between translations in $\mathcal{R}(\pi_v)$ and $T_v S^{n-1}$.*

(1) $\langle \nabla(f \circ \pi_v^{-1})(\pi_v(u)), x \rangle = \langle v, u \rangle \langle P_{u^\perp} \nabla f(u), \xi \rangle$

(2) $x^T [\mathcal{H}(f \circ \pi_v^{-1})(\pi_v(u))] y$
$= \langle v, u \rangle^2 \xi^T [P_{u^\perp} \mathcal{H} f(u) P_{u^\perp}^T - u[P_{u^\perp} \nabla f(u)]^T - P_{u^\perp} \nabla f(u) u^T - \langle \nabla f(u), u \rangle P_{u^\perp}] \zeta \ .$

Proof. For simplicity of notation, we will denote $\pi_v(u) \in T_v S^{n-1}$ by w.

We note that ξ can be expanded in the basis $p_1, \ldots, p_{n-1}$ of $v^\perp$ via $\langle \xi, p_i \rangle = x_i$, and also that $\xi \perp v$. It follows using Lemma 36 that:

$$\begin{aligned}
\langle \nabla(f \circ \pi_v^{-1})(w), x \rangle &= \sum_{i=1}^{n-1} x_i \partial_i (f \circ \pi_v^{-1})(w) \\
&= \sum_{i=1}^{n-1} x_i \left[\sum_{k=1}^{n-1} \partial_{p_k} f(u) \left[\frac{\delta_{ik}}{\|(1,\ w^T)\|} - \frac{w_i w_k}{\|(1,\ w^T)\|^3} \right] - \frac{\partial_v f(u) w_i}{\|(1,\ w^T)\|^3} \right] \\
&= \frac{\langle \nabla f(u), \xi \rangle}{\|\mathrm{rep}(w)\|} - \frac{\langle \mathrm{rep}(w) \langle \nabla f(u), \mathrm{rep}(w) - v \rangle, \xi \rangle}{\|\mathrm{rep}(w)\|^3} - \frac{\langle \mathrm{rep}(w) \langle \nabla f(u), v \rangle, \xi \rangle}{\|\mathrm{rep}(w)\|^3} \ .
\end{aligned}$$

To see that last equality, we note that $\|\mathrm{rep}(w)\| = \|\sum_{i=1}^n w_i p_i + v\| = \|(1,\ w^T)\|$, and we use that $\xi \perp v$ to get that $\langle \mathrm{rep}(w), \xi \rangle = \sum_{i=1}^{n-1} \langle \xi, p_i \rangle w_i = \sum_{i=1}^{n-1} x_i w_i$ and $\langle \nabla f(u), \xi \rangle = \sum_{i=1}^{n-1} \partial_{p_k} f(u) x_i$ hold.

Since $w = \pi_v(u)$, we have that $\frac{\mathrm{rep}(w)}{\|\mathrm{rep}(w)\|} = u$ and $\|\mathrm{rep}(w)\| = \frac{1}{\langle u, v \rangle}$. It follows that

$$\langle \nabla(f \circ \pi_v^{-1})(w), x \rangle = \langle u, v \rangle \langle \nabla f(u) - u \langle \nabla f(u), u \rangle, \xi \rangle \ .$$

which is equivalent to part (1) of this Lemma.

In showing part (2), we somewhat compress the number of steps to save space, but reasoning of a similar nature is used—i.e., we will once again use the expansions $\langle \xi, p_i \rangle = x_i$ and $\langle \zeta, p_i \rangle = y_i$ as well as the facts that $\xi \perp v$ and $\zeta \perp v$ in a similar fashion as it was used before. Using these properties, we expand from (26) (with some rearrangement of terms) to get:

$$\begin{aligned}
&x^T [\mathcal{H}(f \circ \pi_v^{-1})(w)] y^T \\
&= \|\mathrm{rep}(w)\|^{-2} \xi^T \Big[\mathcal{H}f(u) - \|\mathrm{rep}(w)\|^{-2} [\mathrm{rep}(w)(\mathrm{rep}(w) - v)^T \mathcal{H}f(u) \\
&\quad + \mathcal{H}f(u)(\mathrm{rep}(w) - v)(\mathrm{rep}(w))^T + \mathrm{rep}(w) v^T \mathcal{H}f(u) + \mathcal{H}f(u) v\ \mathrm{rep}(w)^T] \\
&\quad + \|\mathrm{rep}(w)\|^{-4} [\mathrm{rep}(w)(\mathrm{rep}(w - v)^T \mathcal{H}f(u)(\mathrm{rep}(w) - v)\, \mathrm{rep}(w)^T \\
&\quad + \mathrm{rep}(w) v^T \mathcal{H}f(u)(\mathrm{rep}(w) - v)\, \mathrm{rep}(w)^T + \mathrm{rep}(w)(\mathrm{rep}(w - v)^T \mathcal{H}f(u) v\ \mathrm{rep}(w)^T \\
&\quad + \mathrm{rep}(w) v^T \mathcal{H}f(u) v\ \mathrm{rep}(w)^T] \\
&\quad + 3\|\mathrm{rep}(w)\|^{-3}\, \mathrm{rep}(w)\, \mathrm{rep}(w)^T [\langle \nabla f(u), \mathrm{rep}(w) - v \rangle + \langle \nabla f(u), v \rangle] \\
&\quad - \|\mathrm{rep}(w)\|^{-1} [\mathrm{rep}(w) \nabla f(u)^T - \nabla f(u)\, \mathrm{rep}(w)^T + I \langle \nabla f(u), \mathrm{rep}(w) - v \rangle + I \langle \nabla f(u), v \rangle] \Big] \zeta
\end{aligned}$$

We note that $u = \frac{\mathrm{rep}(w)}{\|\mathrm{rep}(w)\|}$. So, after combining related terms, doing the appropriate replacement of u, and noting that $\|\mathrm{rep}(w)\| = \frac{1}{\langle u, v \rangle}$, we get:

$$\begin{aligned}
&x^T [\mathcal{H}(f \circ \pi_v^{-1})(w)] y^T \\
&= \langle u, v \rangle^2 \xi^T \Big[\mathcal{H}f(u) - u u^T \mathcal{H}f(u) - \mathcal{H}f(u) u u^T + u u^T \mathcal{H}f(u) u u^T \\
&\qquad + 3 \langle \nabla f(u), u \rangle u u^T - u \nabla f(u)^T - \nabla f(u) u^T - \langle \nabla f(u), u \rangle I \Big] \zeta \\
&= \langle u, v \rangle^2 \xi^T [(I - u u^T) \mathcal{H}f(u)(I - u u^T)^T - u (P_u \nabla f(u) - \nabla f(u))^T \\
&\qquad - (P_u \nabla f(u) - \nabla f(u)) u^T - \langle \nabla f(u), u \rangle (I - u u^T)] \zeta \\
&= \langle u, v \rangle^2 \xi^T [P_{u^\perp} \mathcal{H}(f(u)) P_{u^\perp}^T - u (P_{u^\perp} \nabla f(u))^T - P_{u^\perp} \nabla f(u) u^T - \langle \nabla f(u), u \rangle P_{u^\perp}] \zeta \ . \qquad \square
\end{aligned}$$

Corollary 39. *Suppose that $x, y \in \mathcal{R}(\pi_v)$, and let $\xi = \mathrm{rep}(x) - v$ and $\zeta = \mathrm{rep}(y) - v$ be the corresponding translations in $T_v S^{n-1}$. Then, we have the following operator equivalences between $\mathcal{R}(\pi_v)$ and $T_v S^{n-1}$ at the tangent point v.*

$$\langle \nabla(f \circ \pi_v^{-1})(\pi_v(v)), x \rangle = \langle \nabla f(v), \xi \rangle \ . \tag{1}$$

$$x^T \mathcal{H}(f \circ \pi_v^{-1})(\pi_v(v)) x = \xi^T [\mathcal{H} f(v) - \langle \nabla f(v), v \rangle I] \zeta \ . \tag{2}$$

Proof. This follows from Lemma 38 by using that $\xi \perp v$ and $\zeta \perp v$. □

We thus write that $\nabla(f \circ \pi_v^{-1})(\pi_v(v)) \simeq \nabla f$ and $\mathcal{H}(f \circ \pi_v^{-1})(\pi_v(v)) \simeq \mathcal{H} f(v) - \langle \nabla f(v), v \rangle I$ since they provide representatives of our operators from $\mathcal{R}(\pi_v)$ in the space $v^\perp$, i.e., the space of displacements from v on $T_v S^{n-1}$. Further, it will suffice when analyzing the resulting Newton-like update to perform analysis in the space $\mathcal{R}(\pi_v)$ with the operators $\nabla(f \circ \pi_v^{-1})$ and $\mathcal{H}(f \circ \pi_v^{-1})$ for which we have derived explicit formulas.